An Introduction to ECharts: The Concise User Manual

Gregory W. Bond
AT&T Labs – Research

Version 1.1 beta
August 18, 2006

The latest version of this document is available from:
http://echarts.org

Abstract

This document provides a concise example-driven overview of the ECharts language and its features. Many simple examples are provided to familiarize the reader with the language's textual and graphical syntax, and its semantics. The document also provides instructions for compiling, executing, debugging and generating documentation for ECharts machines.

Contents

1 Introduction 4
1.1 History . . . 5
1.2 Why Use ECharts? . . . 5

2 Developing an ECharts Machine 5
2.1 Hello World! . . . 5

3 What Can ECharts Do? 8
3.1 Hierarchical Machines . . . 8
3.1.1 Nested Inner Machines . . . 8
3.1.2 Nested External Machines . . . 9
3.1.3 Reflective Invocation . . . 9
3.2 Actions . . . 10
3.2.1 Transition Actions . . . 10
3.2.2 Entry and Exit Actions . . . 10

3.2.3 Action Blocks . . . 13
3.3 Multi-Level Transitions . . . 13
3.4 Or-Machines . . . 14
3.5 And-Machines . . . 14
3.6 Fork and Join Transitions . . . 15
3.7 Transition Guards . . . 17
3.8 Machine Constructors . . . 19
3.9 Receiving a Message . . . 21
3.10 Non-Blocking Execution . . . 26
3.11 '*' Transitions . . . 29
3.12 Sending a Message . . . 31
3.13 External Ports . . . 32
3.13.1 Output Handlers . . . 33
3.13.2 Remote Sends . . . 34
3.14 Machine Arrays . . . 34
3.15 Pseudostates . . . 38
3.15.1 DEFAULT_INITIAL . . . 38
3.15.2 TERMINAL . . . 39
3.15.3 DEEP_HISTORY . . . 41
3.15.4 NEW . . . 43
3.15.5 Summary . . . 44
3.16 Timed Transitions . . . 44
3.17 Submachine Access . . . 47
3.18 Partial State References . . . 49
3.19 Host Language Interface . . . 52

4 The Runtime Model **53**
4.1 Machine Execution . . . 53
4.2 Transition Scheduling . . . 54
4.3 Port Priorities . . . 55
4.4 Transition Priorities . . . 57
4.4.1 Message Class Rule . . . 57
4.4.2 Source Coverage Rule . . . 59
4.4.3 Transition Depth Rule . . . 63
4.5 Message Dequeuing . . . 64
4.5.1 Explicit Message Consumption . . . 65
4.5.2 Implicit Message Deferral . . . 66
4.6 Machine Creation and Destruction . . . 68
4.6.1 Machine Creation . . . 68
4.6.2 Machine Destruction . . . 71
4.7 Shared Data . . . 73
4.7.1 `shared` Machine Modifier . . . 73
4.7.2 Port Variables . . . 76
4.8 Machine and State Access Modifiers . . . 76

5 The Machine Runtime 77
5.1 Initialization . . . 77
5.2 Properties File . . . 78
5.3 Startup Messages . . . 78
5.4 Transition Timer Manager . . . 78
5.5 Monitoring and Logging . . . 78
5.5.1 Monitors . . . 79
5.5.2 Events . . . 79
5.5.3 putEvent() . . . 80
5.5.4 Event Filters . . . 80
5.5.5 Formatters and Logging . . . 81
5.6 Debugging . . . 81
5.6.1 Output Format . . . 82
5.6.2 Message Properties . . . 85
5.6.3 Options . . . 86
5.6.4 Global State Output . . . 87
5.7 Options Summary . . . 88

6 Generating Diagrams 88
6.1 Page Diagrams . . . 89
6.2 Embedded Diagrams . . . 91
6.3 Customizing the Layout . . . 91
6.3.1 Overriding `dot` Layout . . . 92
6.3.2 Overriding `ech2dot` Layout . . . 93
6.4 `dot` Layout Bugs . . . 94

7 Generating Documentation 95
7.1 Interacting with Diagrams . . . 95
7.2 `ech2doc` . . . 96
7.3 `ech2javadoc` . . . 96
7.4 SVG Viewers . . . 99

8 Command Reference 102
8.1 Machine Dependencies . . . 102
8.2 ECHARTSPATH . . . 102
8.3 `ech2java` . . . 103
8.4 `ech2dot` . . . 103
8.5 `ech2doc` . . . 104
8.6 `ech2javadoc` . . . 106
8.7 `javadocpp` . . . 106

9 Coming Soon! 107
9.1 Unit and Regression Testing 107
9.2 Indexed Machine Array Access 107
9.3 Internal Ports 107
9.4 Past and Future State References 107
9.5 Exception Handling 108

10 Coming Later! 108
10.1 Machine Inheritance 108
10.2 Machine Variables 108

References 108

A Building and Using ECharts 109
A.1 Software Requirements 109
A.2 Building ECharts 110
A.3 Using ECharts 110
A.3.1 Direct Invocation 110
A.3.2 Indirect Invocation 111

B Licenses 111
B.1 Common Public License v1.0 111
B.2 ANTLR 2 License 116

1 Introduction

What is ECharts? ECharts is a state machine-based programming language for event-driven systems derived from the standardized UML Statecharts language [6]. ECharts distinguishes itself from other Statecharts dialects by focusing on implementation issues such as determinism and code re-use [2]. Like Statecharts, ECharts supports hierarchical state machines, concurrent machines and a graphical syntax. Unlike Statecharts, ECharts supports a simple textual syntax, machine reuse, multiple transition priority levels to minimize non-determinism, machine arrays, and a new approach to inter- and intra- machine communication. ECharts is a hosted language which means that it is dependent on an underlying programming language such as Java[1]. ECharts has a proven track-record in a large-scale commercial deployment. ECharts is available as open source under the Common Public License Version 1.0 (see Appendix B). Take a look at what ECharts has to offer!

[1] Java is a trademark or registered trademark of Sun Microsystems, Inc. in the United States and other countries. echarts.org is independent of Sun Microsystems, Inc.

1.1 History

ECharts was originally developed at AT&T Labs–Research to support an advanced telecommunications project which evolved into a nationwide commercial product offering [3]. Currently ECharts is used in a number of new research projects at AT&T.

1.2 Why Use ECharts?

An ECharts machine is a specification of event-driven behavior. Application domains that will benefit from using ECharts are telecommunications, web applications, user interfaces, ... really any application domain where a system responds to events occurring in its environment.

ECharts isn't meant to be a stand-alone programming language. ECharts specifies a program's control flow in response to events received from the environment. To specify data operations, ECharts machines rely on host language statements embedded in the machine, for example, Java. Furthermore, an ECharts machine only constitutes a program unit in a larger host language program, it isn't necessarily a program's locus of control.

2 Developing an ECharts Machine

The ECharts SDK (System Development Kit) is split into two major components: (1) a translator and (2) a runtime system. The translator translates an ECharts machine into a host language program module, for example, Java. The runtime system is a host language library comprising the ECharts machine interpreter, that is, the code that actually executes the machine. The runtime system library is linked with the compiled ECharts machine modules and other compiled host language modules to form an executable program. For information on building the ECharts SDK see Appendix A.

2.1 Hello World!

Here's an example of how to program, compile, and run the canonical "Hello World!" program as an ECharts machine. In this example and all subsequent examples, we employ Java as the ECharts host language.

First, create the ECharts machine itself in a file *examples/Example0001.ech*. If you feel like cheating a bit, you'll find this file and all other example files mentioned in this document in the *runtime/java/src/examples* directory of the ECharts SDK.

```
package examples;

public machine Example0001 {
    initial state S1;
    state S2;
    transition S1 - / System.out.println("Hello World!") -> S2;
}
```

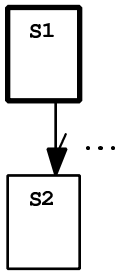

Example0001

For all examples in this document we depict a program using the ECharts textual syntax and its graphical syntax. For a discussion of the many ways that exist to generate machine diagrams like the one above see Section 6.

The `Example0001` machine does nothing more than transition from state `S1`, declared as the machine's `initial` state, to state `S2` executing the action to print "Hello World!" in the process. The graphical syntax depicts the initial state with a bold outline. In ECharts, transition actions are defined following the `/` (forward slash) character. Here the action is a simple expression, `System.out.println("Hello World!")`, specified using Java syntax. In general, Java expression syntax can be used to specify machine actions regardless of the host language. These are translated to the appropriate host language expressions (in this case, Java expressions) appropriately. ECharts also permits otherwise unacceptable expressions to be wrapped in host code delimiters `<* *>`. Further discussion concerning the ECharts/host language interface can be found in Section 3.19.

To translate the machine to the host language, invoke the ECharts translator (provided with the ECharts SDK) from the *examples* directory. In this example and subsequent examples, we show a command line prompt as a percent character `%`.

```
% ech2java --echartspath .. Example0001.ech
```

For information on configuring your platform to run ECharts commands like `ech2java` see Appendix A. This command will create a

file *examples/Example0001.java*. Note the use of the `--echartspath` command line option. This option informs the translator where to locate the `examples` package declared in our machine. It is analogous to Java's `javac -classpath` command line option. See Section 8 for more `ech2java` command line options.

To compile and link the translated machine, invoke the host language's compiler from the *examples* directory:

```
% javac -classpath ..:../../echarts.jar Example0001.java
```

This command creates the file *examples/Example0001.class*. Note that the `-classpath` option references the *echarts.jar* file containing the ECharts Java runtime system (part of the ECharts SDK build).

The next step is to write a host language program whose mainline procedure calls the machine. To do this create a file *examples/HelloWorld.java*:

```
package examples;

public class HelloWorld {
    public static final void main(String argv[]) {
        try {
            new Example0001().run();
        } catch (Exception e) {
            e.printStackTrace();
        }
    }
}
```

All this program does is create a new instance of an Example0001 machine and then invoke its `run()` method. The `run()` method returns after the ECharts machine runtime has determined that it can execute no further machine transitions.

Now, from the *examples* directory, compile and link the mainline program:

```
% javac -classpath ..:../echarts.jar HelloWorld.java
```

This will produce the file *examples/HelloWorld.class*. As the final step, run the program:

```
% java -classpath ..:../echarts.jar examples.HelloWorld
Hello World!
```

There you have it. Naturally, the steps described above for building an executable program can be automated using your favorite build tool such as ant or make. A GNU Makefile and and Apache Ant build.xml file are included in ECharts SDK for building the example machines in this document. For more information about ECharts runtimes and translators see Sections 5 and 8, respectively.

3 What Can ECharts Do?

Now that you've seen how to write and run a simple ECharts program, we're ready to take a look a closer look at the ECharts language and what it can do. In the following sub-sections we provide a brief overview of ECharts language features.

3.1 Hierarchical Machines

ECharts permits nesting a machine in a machine's state. ECharts supports two approaches to nesting machines: (1) inner machines and (2) external machines.

3.1.1 Nested Inner Machines

Here's a simple example of a nested inner machine:

```
package examples;

public machine Example0002 {
    initial state S1: {
        initial state S1_1;
        state S1_2;
        transition S1_1 - /
            System.out.println("Hello World")
        -> S1_2;
    }
}
```

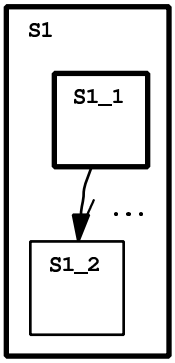

Example0002

In this example, an inner machine is nested in the initial state `S1` of the root (or top-level) machine, `Example0002`. Variables and methods defined in the parent machine are accessible to nested inner machines.

3.1.2 Nested External Machines

ECharts also supports nesting parameterized externally defined machines. This, in turn, supports machine re-use. Here's an example:

```
package examples;

public machine Example0003 {
    initial state S1: Example0001();
}
```

Example0003

In this example, the Example0001 machine is nested in state `S1` of the Example0003 machine. We discuss parameterized nested machines in Section 3.8. Access permissions for external machines is discussed in Section 4.8.

3.1.3 Reflective Invocation

In the previous example, the nested machine's identity is declared as a constant. Using ECharts reflective invocation, the nested machine

can also be declared as a variable whose value is the machine's fully qualified class name as shown in this example:

```
package examples;

public machine Example0004 {

    <* static final String machineName =
       "examples.Example0001" *>
    <* static final Object[] machineArguments = null *>

    initial state S1: reflect(machineName, machineArguments);
}
```

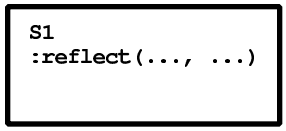

Example0004

This example also shows how host language statements are wrapped in host language delimiters `<* *>`. When Java is the host language, variable declarations, method declarations and inner class declarations must be wrapped in host language delimiters. For more details concerning the ECharts/host language interface see Section 3.19.

3.2 Actions

There are four kinds of action that may execute during a machine's execution: (1) a transition action, (2) an entry action, (3) an exit action, (4) a constructor.

3.2.1 Transition Actions

As we've seen in the examples above, an ECharts machine can perform actions when its machine transitions fire.

3.2.2 Entry and Exit Actions

A machine can also perform actions upon entry to or exit from a machine's state. A state's exit action is executed when a firing transition explicitly references the state as a source state. A state's entry action is executed when a firing transition explicitly references the state as a target state. For example:

```
package examples;

public machine Example0005 {
    initial state S1
        entry System.out.println("Entering S1")
        exit System.out.println("Exiting S1");
    state S2
        entry System.out.println("Entering S2")
        exit System.out.println("Exiting S2") : {
        initial state S2_1
            entry System.out.println("Entering S2_1")
            exit System.out.println("Exiting S2_1");
    }
    state S3;
    transition S1 - / System.out.println("Hello") -> S2;
    transition S2 - / System.out.println("World!") -> S3;
}
```

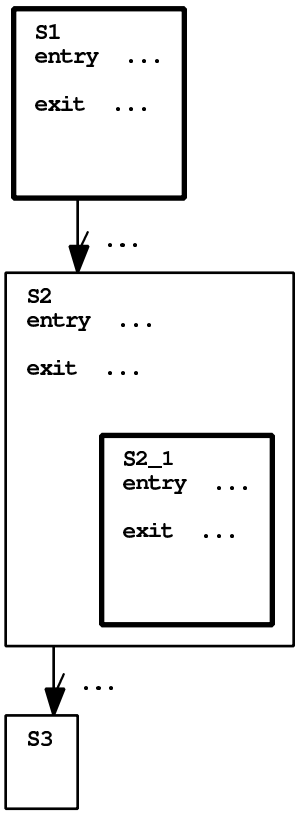

Example0005

This program produces the following output:

```
Exiting S1
Hello
Entering S2
Exiting S2
World!
```

When the first transition fires, its source explicitly references state `S1` which results in its exit action being executed. The same transition's target explicitly references state `S2` which results in its entry action being executed. When the second transition fires, its source explicitly references state `S2` which results in its exit action executing. Notice that the entry action for `S1` and the entry and exit actions for `S2_1` are not executed because neither of these two states are explicitly referenced by transition source or target. Entry actions execute top-down as a new machine state is entered. Exit actions execute

bottom-up as an old machine state is exited. More details concerning entry and exit actions are discussed in the section devoted to machine pseudostates in Section 3.15.

3.2.3 Action Blocks

In the previous example, only single line action statements were uses. Actions can also be grouped together into multi-line action blocks as follows:

```
package examples;

public machine Example0006 {
    initial state S1
        exit {
          System.out.print("Exiting ");
          System.out.print("from ");
          System.out.println("S1");
        }
    state S2;
    transition S1 --> S2;
}
```

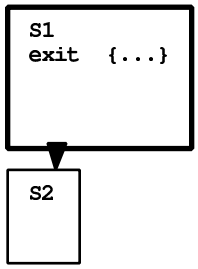

Example0006

As you can see from this example, action blocks are enclosed by curly brackets { } and individual actions are delimited with a semicolon.

3.3 Multi-Level Transitions

In the previous examples, transitions have referenced states in the same machines the transitions are declared in. ECharts also supports transitions referencing states declared in submachines, as shown in this example:

```
package examples;

public machine Example0007 {

    initial state S1: {
        initial state S1_1;
    }
    state S2: {
        initial state S2_1;
    }
    transition S1.S1_1 --> S2.S2_1;
}
```

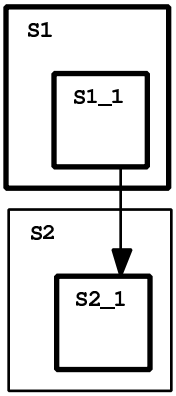

Example0007

Here, the transition source references state `S1_1` of the submachine nested in `S1`, and the transition target references state `S2_1` of the submachine nested in state `S2`.

Machine state access permissions can be specified to dictate whether a state can be referenced by a transition. See Section 4.8 for further discussion.

3.4 Or-Machines

So far, all our example machines have been instances of *or-machines*. An or-machine corresponds to the traditional notion of a state machine in that it has exactly one current state.

3.5 And-Machines

Unlike an or-machine, all states of an *and-machine* are current states. As such, and-machines support concurrent machine operation. Here's

an example:

```
package examples;

public concurrent machine Example0008 {
    state S1: Example0001();
    state S2: Example0001();
}
```

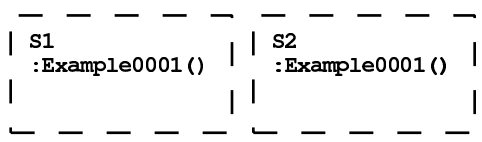

Example0008

This machine concurrently runs two instances of the Example0001 machine. Peer concurrent states are graphically depicted with dashed lines. Here's the machine's output:

```
Hello World!
Hello World!
```

The machine declaration includes the `concurrent` machine modifier, marking it as an and-machine. Inner submachines may also be declared as concurrent. Since both of the machine's states are current states there is no need to declare a state with the `initial` state modifier. For more insight into how concurrently executing machines operate see Section 4.2. Also, in Section 4.7 we provide details concerning data sharing amongst concurrent states. In Sections 4.4.1 and 4.7.1 we discuss in detail how transitions are chosen to fire across concurrent submachines.

3.6 Fork and Join Transitions

When an and-machine is used, the programmer may want a transition to reference more than one concurrent state. To do this ECharts supports the concept of fork and join transitions. Here's an example:

```
package examples;

public machine Example0009 {
    initial state S1: Example0008();
    state S2: Example0008();
    transition [ S1.S1.S2, S1.S2.S2 ] - /
```

```
        System.out.println("Hola Mundo!") ->
    [ S2.S1.S2, S2.S2.S2 ];
}
```

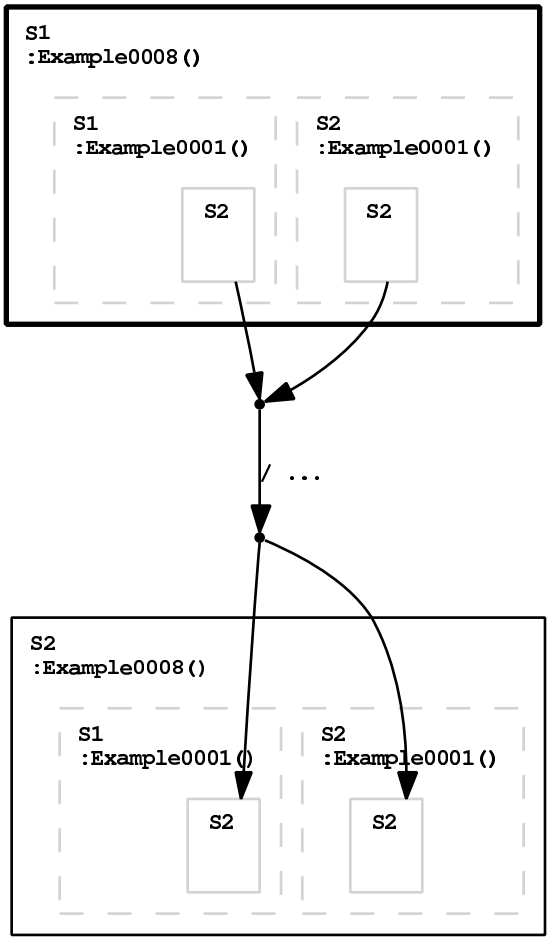

Example0009

The graphical syntax depicts fork and join points with small filled circles. States of external submachines that are explicitly referenced by a machine's transitions are graphically depicted with gray borders. If an external machine's states aren't explicitly referenced then they are not shown in order to reduce clutter.

Here's the machine's output:

```
Hello World!
Hello World!
```

```
Hola Mundo!
Hello World!
Hello World!
```

The source states referenced by the (join) transition in the root machine `Example0009` are sub-states of of the and-machine `Example0008` nested in `S1`. Similarly the target states referenced by the (fork) transition are sub-states of the and-machine `Example0008` nested in `S2`.

The previous example illustrated the use of fork/join transitions in an or-machine. Transitions can also be defined for and-machines as shown in this example:

```
package examples;

public concurrent machine Example0010 {
     state S1: {
         initial state S1_1;
         state S1_2;
     }
     state S2: {
         initial state S2_1;
         state S2_2;
     }
     transition [ S1.S1_1, S2.S2_1 ] --> [ S1.S1_2, S2.S2_2 ];
}
```

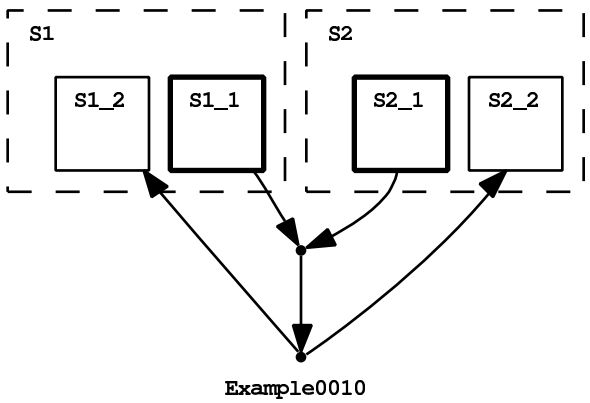

We discuss and-machine transitions further in Section 3.18.

3.7 Transition Guards

A predicate, called a guard, can be specified for a transition. In order for a transition to fire it is necessary for the transition's guard to

evaluate to true. For example:

```
package examples;

public machine Example0011 {
    <* static final int i = 0 *>
    initial state S1;
    state S2;
    state S3;
    transition S1 - [ i == 0 ] /
        System.out.println("i == 0") -> S2;
    transition S1 - [ i > 0 ] /
        System.out.println("i > 0") -> S3;
}
```

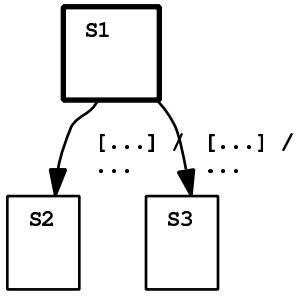

Example0011

Because `i` is initialized to `0`, only the first transition will be enabled so this machine will transition to state `S2`.

Guards can also be grouped into if-then-elif-else style compound statements as shown in the following example:

```
package examples;

public machine Example0045 {
    <* private int iterations = 0 *>
    <* static final int MAX_ITERATIONS = 2 *>
    initial state S1;
    state S2;
    transition S1 -
        [ iterations < MAX_ITERATIONS ] {
            [ iterations == 0 ] / {
                System.out.println("iteration: 1");
                iterations++;
```

```
            } -> S1
            else [ iterations == 1 ] / {
                System.out.println("iteration: 2");
                iterations++;
            } -> S1
        }
        else -> S2;
}
```

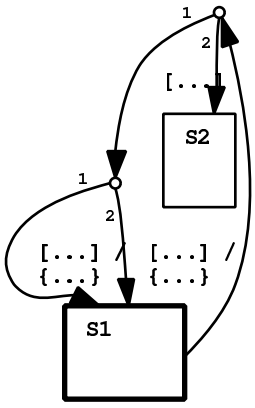

Example0045

The graphical syntax depicts branching points as small hollow circles and branch segments include branch numbers to indicate guard evaluation order. And here's the output from the example:

```
iteration: 1
iteration: 2
```

In this example, there is only a single transition with nested compound targets. The machine iterates two times from state `S1` to state `S1`, each time printing out the iteration number. When the maximum number of iterations is reached, the machine transitions to state `S2`. With compound targets, the guards are evaluated in the usual order as, for example, in Java. We discuss guards further in Section 3.9.

3.8 Machine Constructors

ECharts machines may declare constructors if desired. By default, all machines implicitly declare a zero-argument constructor. This has been the case for all the examples to this point. The next example shows how to declare a machine constructor and how to invoke the constructor from another machine.

```
package examples;

public machine Example0017 {
    <* private String outString = null *>
    initial state S1;
    state S2: Example0018(outString);
    transition S1 - / outString = "Hello World!" -> S2;
}
```

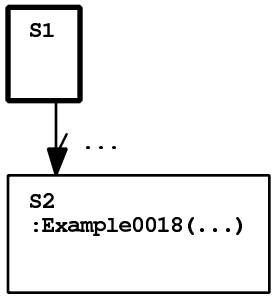

Example0017

```
package examples;

public machine Example0018 {

    <* final private String outString *>

    public Example0018(String outString) {
        this.outString = outString;
    }
    initial state S1;
    state S2;
    transition S1 - / System.out.println(outString) -> S2;
}
```

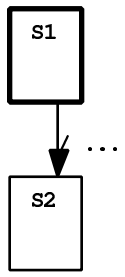

Example0018

The `Example0017` machine invokes the `Example0018` machine when it transitions from `S1` to `S2`. When the transition fires the transition action initializes the value of the `Example0018` machine constructor parameter. The parameter value is referenced by the `Example0018` machine constructor when its own transition fires, resulting in the following output:

```
Hello World!
```

3.9 Receiving a Message

All the transitions in the examples discussed so far have been *messageless* transitions. We now introduce *message* transitions. An ECharts machine reacts to its environment when it receives messages on external ports specified by its message transitions. Consider the following example machine:

```
package examples;

public machine Example0012 {

    <* final private ExternalPort p1 *>

    public Example0012(ExternalPort p1) {
        this.p1 = p1;
    }
    initial state S1;
    state S2;
    transition S1 - p1 ? String /
        System.out.println("string: " + message)
    -> S2;
    transition S1 - p1 ? Float /
        System.out.println("float: " + message)
```

```
    -> S2;
}
```

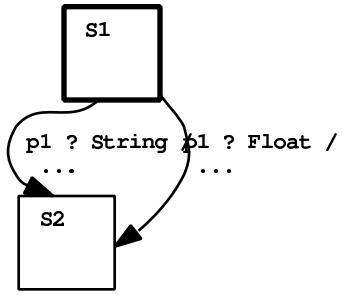

Example0012

Here's the machine's environment:

```
package examples;

import org.echarts.ExternalPort;
import org.echarts.MachineThread;

public class Example0012Environment {

    final static private ExternalPort p1 =
        new ExternalPort("p1");

    final static public void main(String[] argv) {
        try {
            new MachineThread(new Example0012(p1)).start();
            p1.input("Hello World!");
        } catch (Exception e) { e.printStackTrace(); }
    }
}
```

First, the machine's environment creates an instance of the machine and runs it on a separate thread. The `MachineThread` class is a convenience class that does nothing more than call the machine's `run()` method. The machine's thread blocks waiting for a message to arrive on `p1`. Then the machine's environment puts a message, a `String` instance, in external port `p1`'s input queue. The arrival of the message unblocks the machine's thread, and the transition specifying a `String` message fires, printing out the message. ECharts also supports non-blocking machine execution, described in Section 3.10.

A message transition uses the port receive syntax `p1 ? Classname` to specify waiting on external port `p1` for a message that is an instance of the `Classname` class. When the transition's guard is evaluated and when the transition's action is executed, the message instance is made available as a distinguished variable named `message`. The port instance is made available as a distinguished variable named `port`. While this variable's value is redundant here, it is useful for '*' transitions (see Section 3.11).

When enabled message transitions specify the same port and same received message, then the ECharts runtime will give highest priority to the transition specifying the most specific message class. Here's an example:

```
package examples;

public machine Example0047 {

    <* final private ExternalPort p1 *>

    public Example0047(ExternalPort p1) {
        this.p1 = p1;
    }

    initial state S1;
    state S2;
    transition S1 - p1 ? Object /
        System.out.println("Object: " + message)
    -> S2;
    transition S1 - p1 ? String /
        System.out.println("String: " + message)
    -> S2;
}
```

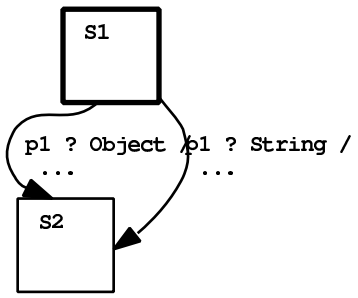

Example0047

```
package examples;

import org.echarts.ExternalPort;

public class Example0047Environment {

    public static final void main(String[] argv) {
        try {
            final ExternalPort p1 = new ExternalPort("p1");
            p1.input("Hello World!");
            new Example0047(p1).run();
        } catch (Exception e) { e.printStackTrace(); }
    }
}
```

This example shows the machine's `run()` method being invoked directly by the environment's `main()` method as an alternative to creating a separate `MachineThread` instance as shown in the previous example. When the machine is executed the runtime chooses the second transition to fire because it specifies the `String` class which is more specific than the `Object` class (which is a superclass of the `String` class). This transition priority rule is the first of three transition priority rules supported by the ECharts runtime. These rules are discussed in more detail in Section 4.4.

A given external port can be referenced by more than one transition in a machine, for example, by transitions in different concurrent states of an and-machine:

```
package examples;

public concurrent machine Example0013 {

    <* final private ExternalPort p1 *>

    public Example0013(ExternalPort p1) {
        this.p1 = p1;
    }
    state S1: {
        initial state S1_1;
        state S1_2;
        transition S1_1 - p1 ? String /
            System.out.println(message)
```

```
        -> S1_2;
    }
    state S2: {
        initial state S2_1;
        state S2_2;
        transition S2_1 - p1 ? String /
            System.out.println(message)
        -> S2_2;
    }
}
```

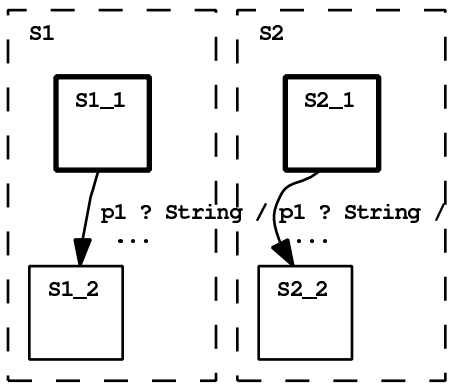

Example0013

Here's the previous machine's environment:

```
package examples;

import org.echarts.ExternalPort;
import org.echarts.MachineThread;

public class Example0013Environment {

    public static final void main(String[] argv) {
        try {
            final ExternalPort p1 = new ExternalPort("p1");
            new MachineThread(new Example0013(p1)).start();
            p1.input("Hello");
            p1.input("World!");
        } catch (Exception e) { e.printStackTrace(); }
    }
}
```

In this example, the machine's environment puts two messages in `p1`'s input queue. The two transitions in the machine's concurrent states fire in succession producing one of two possible outputs. Here's one possible output:

```
Hello
World!
```

Normally, a port variable's value will remain constant throughout the life of a machine. However, if a port's value is expected to change then there are constraints on how that value may be changed. This is discussed in more detail in Section 4.7.2.

We discuss the blocking machine execution cycle in more detail in Section 4.1. A non-blocking alternative is discussed in the next section, Section 3.10. We discuss how messages received from the environment are dequeued in more detail in Section 4.5.

3.10 Non-Blocking Execution

In the last two examples we discussed how the thread that invokes a machine's `run()` method can block while awaiting the arrival of messages specified by message transitions in the machine's current state. ECharts also provides a non-blocking `run()` method that can be useful in some application domains such as HTTP or SIP servlet containers. Here's how the machine in the previous example would be executed using a non-blocking approach:

```
package examples;

import org.echarts.ExternalPort;
import org.echarts.Machine;

public class Example0014Environment {

    public static final void main(String[] argv) {
        try {
            final ExternalPort p1 = new ExternalPort("p1");
            Machine machine = new Example0014(p1);
            machine.run(p1, "Hello");
            machine.run(p1, "World!");
        } catch (Exception e) { e.printStackTrace(); }
    }
}
```

Here, the machine's environment creates a machine instance (where `Machine0014` is identical to `Machine0013`) and then invokes the non-blocking `run()` method two times. The arguments for the `run()` method specify the port `p1` and the message associated with the port.

In the case where the non-blocking `run()` method is invoked specifying port p and no active message transitions specify p in the machine's current state then the message will be enqueued on p's input queue in accordance with the implicit message deferral model supported by the ECharts runtime (see Section 4.5.2). Messages enqueued as a result of earlier `run()` invocations may be dequeued and processed by later `run()` invocations in accordance with the runtime's explicit message consumption model (see Section 4.5.1). Here's an example:

```
package examples;

public machine Example0048 {

    <* final private ExternalPort p1 *>
    <* final private ExternalPort p2 *>

    public Example0048(ExternalPort p1, ExternalPort p2) {
        this.p1 = p1;
        this.p2 = p2;
    }
    initial state S1;
    state S2;
    state S3;
    transition S1 - p2 ? String /
        System.out.println(port + ": " + message)
    -> S2;
    transition S2 - p1 ? String /
        System.out.println(port + ": " + message)
    -> S3;
}
```

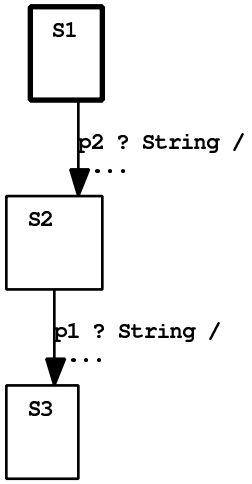

Example0048

```
package examples;

import org.echarts.ExternalPort;

public class Example0048Environment {

    final static public void main(String[] argv) {
        try {
            final ExternalPort p1 = new ExternalPort("p1");
            final ExternalPort p2 = new ExternalPort("p2");
            final Example0048 machine =
                new Example0048(p1, p2);
            machine.run(p1, "Hello");
            machine.run(p2, "World");
        } catch (Exception e) { e.printStackTrace(); }
    }
}
```

And here's the example's output:

```
p2: World
p1: Hello
```

In this example, the environment invokes the machine's non-blocking `run()` method two times: first with a message for port `p1` and then with

a message for port p2. Since the only message transition specifying p1 is not active in the machine's initial state S1, then p1's message is enqueued on p1's input queue and the machine's state remains unchanged the first time run() is invoked. When run() is invoked a second time, the message transition specifying p2 fires. Then the method dequeues p1's message and fires p1's transition because the transition is active in state S2.

A blocking approach to machine execution is discussed in the previous section, Section 3.9.

3.11 '*' Transitions

A variation on the message transition is the '*' ("any port") transition. The '*' transition is a message transition that specifies a wildcard port. In particular, it matches any port that is explicitly mentioned in some other message transition in a machine's current state. Here's an example machine using a '*' transition:

```
package examples;

public machine Example0015 {

    <* final private ExternalPort p1 *>
    <* final private ExternalPort p2 *>

    public Example0015(ExternalPort p1, ExternalPort p2) {
        this.p1 = p1;
        this.p2 = p2;
    }
    initial state S1;
    state S2;
    transition S1 - p1 ? Float /
        System.out.println(message)
    -> S2;
    transition S1 - p2 ? Integer /
        System.out.println(message)
    -> S2;
    transition S1 - * ? String / {
            System.out.println(port);
            System.out.println(message);
        }
    -> S2;
}
```

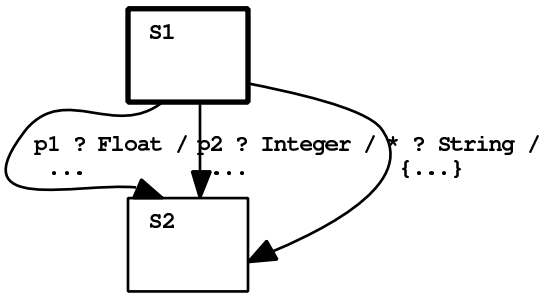

Example0015

And here's the machine's environment:

```
package examples;

import org.echarts.ExternalPort;
import org.echarts.MachineThread;

public class Example0015Environment {

    public static final void main(String[] argv) {
        try {
            final ExternalPort p1 = new ExternalPort("p1");
            final ExternalPort p2 = new ExternalPort("p2");
            new MachineThread(new Example0015(p1, p2)).start();
            p1.input("Hello World!");
        } catch (Exception e) { e.printStackTrace(); }
    }
}
```

In this example, the machine's environment puts a `String` instance in external port `p2`'s input queue. The only transition to explicitly reference `p2` specifies an `Integer` message, however, the '*' transition does specify a `String` instance and, since the '*' matches any port specified by a transition in the current state, it matches `p2` and fires. Here's the machine's output:

```
p1
Hello World!
```

Notice that the '*' transition action references a distinguished variable named `port` whose value is the matched port.

3.12 Sending a Message

Naturally, in addition to receiving messages, a machine can also send messages to its environment via external ports as shown in this example:

```
package examples;

public machine Example0019 {

    <* final private ExternalPort p1 *>

    public Example0019(ExternalPort p1, String message) {
        this.p1 = p1;
        p1 ! message;
    }
    initial state S1;
    state S2;
    transition S1 - p1 ? String /
        System.out.println(message)
    -> S2;
}
```

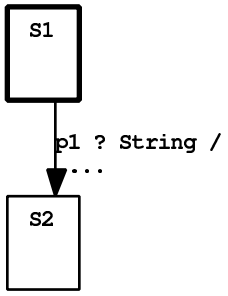

Example0019

```
package examples;

import org.echarts.ExternalPort;
import org.echarts.MachineThread;

public class Example0019Environment {

    public static final void main(String[] argv) {
```

```
        try {
            final ExternalPort p1 = new ExternalPort("p1");
            final ExternalPort p2 = new ExternalPort("p2");
            p1.setPeer(p2);
            p2.setPeer(p1);
            new MachineThread(
                new Example0019(p1, "Hello")).start();
            new MachineThread(
                new Example0019(p2, "World")).start();
        } catch (Exception e) { e.printStackTrace(); }
    }
}
```

In this example, the environment creates two external ports and sets them as each other's peers using the `setPeer()` method. Setting a port's peer means that when a message is sent on a port, the message will be placed in the peer port's input queue. (In Section 3.13 we discuss another way for the environment to receive messages from a machine.) Next, two instances of the `Example0019` machine are created and run on separate threads. Each machine's constructor is invoked with different values for its port and message. Executing the machine constructor sends the specified message on the specified port which enqueues the message on the peer port's input queue. The syntax for a message send operation is `port ! message`. Since port input queue size is unbounded the port send operation is non-blocking. The port send operation can be used anywhere actions are permitted i.e. in state entry/exit actions, transition actions, or machine constructors. Finally, each machine waits for a message to arrive on its own input queue. When the message arrives, its transition fires and the message is printed. One possible output from this example is:

```
World
Hello
```

We discuss how messages received from the environment are dequeued in more detail in Section 4.5.

3.13 External Ports

As we have seen in the previous examples, an external port maintains a input queue of messages that it has received but has not yet processed. To send messages on a port, the programmer must configure the port to output the messages to a particular destination.

3.13.1 Output Handlers

`Example0019` showed that in order to enqueue messages sent out on a port, the port can be peered with another port. However, queuing output messages is not always necessary. Instead, it may be desirable to simply invoke a listener method in the environment when a machine outputs a message. This can be accomplished by overriding the external port's `output()` method as shown in the next example:

```
package examples;

public machine Example0020 {

    <* final private ExternalPort p1 *>

    public Example0020(ExternalPort p1) {
        this.p1 = p1;
    }
    initial state S1;
    state S2;
    transition S1 - / p1 ! "Hello World!" -> S2;
}
```

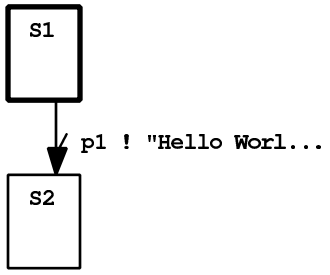

Example0020

```
package examples;

import org.echarts.ExternalPort;

public class Example0020Environment {

    public static final void main(String[] argv) {
        try {
            final ExternalPort p1 = new ExternalPort("p1") {
```

```
                final public void output(final Object message)
                    throws Exception {
                    System.out.println(message);
                }
            };
            new Example0020(p1).run();
        } catch (Exception e) { e.printStackTrace(); }
    }
}
```

In this example the machine's environment creates an instance of an (anonymous) `ExternalPort` subclass. The subclass specifies that the `output()` method should simply print any message sent on that port, which is precisely what happens when the `Example0020` machine fires its transition.

3.13.2 Remote Sends

Messages may also be sent to remote external port instances. For example, imagine that the two `Example0020` machine instances in the previous example are running on different machines. The `ExternalPort` class supports this via its `getRemote()` method. This method returns a remote reference to the port. When a message is input to the remote port reference, either directly by calling the port's `input()` method, or indirectly by invoking a message send operation `!` on the port's peer, the message is placed in the actual port's input queue. For the Java runtime system this is accomplished via Java's RMI.

3.14 Machine Arrays

In addition to the and-machine, ECharts supports another kind of concurrent machine known as the machine array. A machine array is a bounded array of concurrent machines (elements) nested within a single state. Unlike the concurrent states of an and-machine that are all created when the and-machine is created, machine array elements must explicitly be created one-by-one by an ancestor machine at run-time.

Here's an example of machine array element creation:

```
package examples;

public machine Example0021 {

    <* private final String[] messages =
```

```
        new String[]{"Hello", "World"} *>
    <* private final int bound = messages.length *>
    <* private int created = 0 *>

    initial state S1[bound]: Example0018(messages[created - 1]);
    state S2;
    transition S1 - [ created < bound ] / created++ -> S1.NEW;
    transition S1 - [ created == bound ] -> S2;
}
```

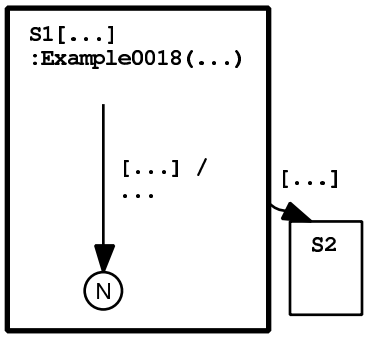

Example0021

In this example, state `S1` declares a machine array of size `bound = 2`: one element for each string element in the `messages` array. Machine array instances in `S1` are declared to be `Example0018` machines. When the `Example0021` machine is initially created it is important to realize that *no* `Example0018` instances are created. A machine instance is created only when the transition from `S1` to `S1.NEW` fires (see Section 6.4 for an explanation of the graphical depiction of this transition). The `NEW` substate referred to by the transition is a special kind of state, called a *pseudostate*, that is defined for all machines. This pseudostate is graphically depicted by the letter 'N' surrounded by a circle. (This pseudostate and others are described in more detail in Section 3.15.) When then `NEW` pseudostate is referred to in the context of a machine array, the effect is to create a new machine instance in the machine array. The states of any other pre-existing machine instances in the array are unaffected. After the `Example0021` machine creates both `Example0018` submachines, it transitions to `S2`. The resulting output is:

```
Hello
World
```

When more than one machine instance exists in a machine array, you may wonder what machine is referenced when a transition refer-

ences the array. There are two simple rules: (1) a transition source state reference is satisfied if *any* machine in the array satisfies the source state; (2) a target state reference refers to *all* pre-existing machines in the array. These rules are highlighted in the following example:

```
package examples;

public machine Example0022 {

    <* private final String[] messages =
        new String[]{"Hello", "World"} *>
    <* private final int bound = messages.length *>
    <* private int created = 0 *>

    initial state S1[bound]: {
        initial nonterminal state S1_1;
        state S1_2;
        state S1_3;
        transition S1_2 - /
            System.out.println(messages[getMachineIndex()])
        -> S1_3;
    }
    transition S1 - [ created < bound ] / created++ -> S1.NEW;
    transition S1.S1_1 - [ created == bound ] -> S1.S1_2;
}
```

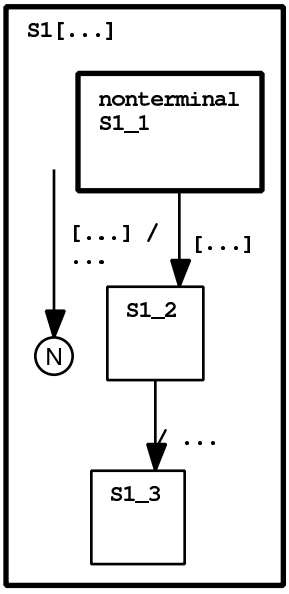

Example0022

Like `Example0021`, this machine first creates two submachines in the machine array declared for state `S1` (see Section 6.4 for an explanation of the graphical depiction of this transition), however, in this case the submachine is defined as an inner machine instead of an external machine. Once created, each submachine instance is prevented from transitioning from their initial state `S1_1`. (The reason for including the `nonterminal` state modifier for `S1_1` will become clear when we discuss machine destruction in Section 4.6.) Once both submachine instances are created the source state `S1.S1_1` referenced by the parent machine's second transition is satisfied (see Section 6.4 for an explanation of the graphical depiction of this transition). This is because there exists at least one submachine in state `S1_1`. In fact, both submachines happen to be in this state. When the transition fires, it updates the state of *both* submachine instances to `S1_2` enabling the submachines' second transition to fire. The `getMachineIndex()` method referenced in the transition's action returns the index of the submachine in the parent machine's machine array. So, when the submachine transitions fire, they print out the submachines' respective messages, one possible output being:

```
Hello
World
```

When a machine array is referenced more than once by a join transition (see Section 3.6) then the interpretation is that all references are

to the same machine array element.

In Section 4.4.1 we discuss how message transitions are chosen to fire across machine array elements. In Section 4.7.1, we discuss how messageless transitions are chosen to fire across machine array elements. In Section 3.17 we discuss how an ancestor machine can access data and methods in machine array elements. In Section 4.7 we discuss data sharing amongst machine array elements. The next version of ECharts will permit individual array elements to be explicitly referenced (see Section 9.2).

3.15 Pseudostates

A small number of pseudostates are defined for all machines. Unlike a real state, a pseudostate represents a machine property. Pseudostates can be referenced in a transition's source or target state, however, a machine's current state cannot be a pseudostate since a pseudostate is not a real machine state by definition.

3.15.1 DEFAULT_INITIAL

The `DEFAULT_INITIAL` pseudostate references a machine's declared initial state. This pseudostate can only be referred to as part of a transition's target state reference. If the machine is an or-machine that has not declared an initial state, then referencing its `DEFAULT_INITIAL` pseudostate results in a translator error. If the machine is an and-machine or a machine array, then referencing its `DEFAULT_INITIAL` pseudostate is a short-hand for referencing the `DEFAULT_INITIAL` pseudostates of each of the machine's concurrent submachines. A reference to a machine's `DEFAULT_INITIAL` pseudostate is treated as an explicit reference to the machine's initial state. Therefore, as discussed in Section 3.2.2, if an entry action is defined for the initial state then it will be executed. While the interpretation of this rule is straightforward for or-machines or machine arrays, it is more subtle for and-machines. If the `DEFAULT_INITIAL` pseudostate of an and-machine is referenced, then any entry actions defined for the and-machine's concurrent states will be executed, however, any entry actions defined for the initial states of the concurrent states' submachines will not be executed. Here's a simple example using the `DEFAULT_INITIAL` pseudostate:

```
package examples;

public machine Example0023 {
    initial state S1;
    state S2: {
```

```
        initial state S2_1
            entry System.out.println("Hola Mundo");
        state S2_2;
        transition S2_1 - /
            System.out.println("Hello World")
        -> S2_2;
    }
    transition S1 --> S2.DEFAULT_INITIAL;
}
```

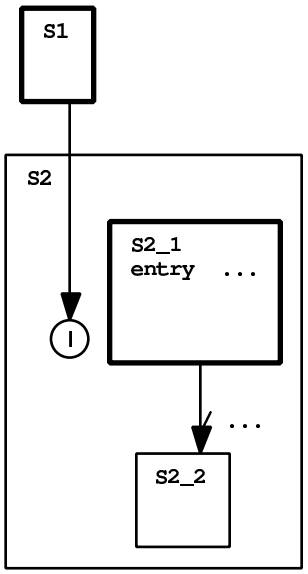

Example0023

When the parent transition fires it references the `DEFAULT_INITIAL` pseudostate of the submachine defined in state `S2`. This pseudostate is graphically depicted by the letter 'I' enclosed by a circle. Since this is treated as an explicit reference to the submachine's initial state, the entry action defined for state `S2_1` is executed followed by the firing of the submachine's transition. The resulting output is:

```
Hola Mundo
Hello World
```

3.15.2 TERMINAL

The `TERMINAL` pseudostate can only be referred to as part of a transition's source state reference. A `TERMINAL` source state reference is satisfied if the referenced machine is in a *terminal* state. An or-machine is

defined to be in a terminal state if the machine's current state is a state with no submachines and no outbound transitions. An and-machine is in a terminal state if all of its concurrent submachines are in terminal states. (We further discuss the notion of a terminal state when we discuss machine destruction in Section 4.6.) If a `TERMINAL` pseudostate references a machine array, then the reference is satisfied if at least one of the machine array elements is in a terminal state. (This is consistent with the rules for transition source state references discussed in Section 3.14.) Similar to what was discussed for the `DEFAULT_INITIAL` pseudostate in Section 3.15.1, a `TERMINAL` pseudostate reference is considered to be an explicit reference to a machine's terminal state. Therefore, any exit actions defined for a terminal state referenced with the `TERMINAL` pseudostate will be executed. Here's a simple example:

```
package examples;

public machine Example0024 {
    initial state S1: {
        initial state S1_1;
        state S1_2
            exit System.out.println("Hola Mundo");
        transition S1_1 - /
            System.out.println("Hello World")
        -> S1_2;
    }
    state S2;
    transition S1.TERMINAL --> S2;
}
```

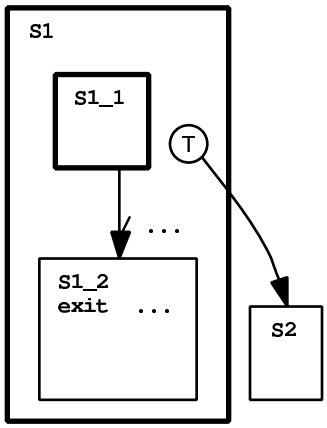

Example0024

The terminal pseudostate is graphically depicted by the letter 'T' enclosed by a circle. The first thing to occur is the transition in state `S1`'s submachine fires. This results in the submachine entering a terminal state `S1_2`. This condition satisfies the source state reference of the parent machine's transition `S1.TERMINAL` causing the parent to transition to state `S2`. Since the reference to the `TERMINAL` pseudostate is treated as an explicit reference to `S1_2`, then `S1_2`'s exit action is executed when the parent transition fires. Here's the machine's output:

```
Hello World
Hola Mundo
```

3.15.3 DEEP_HISTORY

The `DEEP_HISTORY` pseudostate can only be referred to as part of a transition's target state reference. `DEEP_HISTORY` references a machine's current state and the current states of any of its submachines. Unlike references to the `DEFAULT_INITIAL` pseudostate, or the `TERMINAL` pseudostate discussed in Sections 3.15.1 and 3.15.2, respectively, a reference to a machine's `DEEP_HISTORY` pseudostate is *not* considered to be an explicit reference to the machine's current state. Therefore, any entry actions defined for the machine's current state will *not* be executed when referenced via the machine's `DEEP_HISTORY` pseudostate. Here's an example that should clarify this:

```
package examples;

public machine Example0025 {

    <* final private ExternalPort p1 *>

    public Example0025(ExternalPort p1) {
        this.p1 = p1;
    }
    initial state S1: {
        initial state S1_1;
        state S1_2
            entry System.out.println("Hola Mundo");
        state S1_3;
        transition S1_1 - /
            System.out.println("Hello World")
        -> S1_2;
        transition S1_2 - p1 ? Integer /
            System.out.println(message)
```

```
        -> S1_3;
    }
    state S2;
    transition S1 - p1 ? String /
        System.out.println(message)
    -> S1.DEEP_HISTORY;
    transition S1.TERMINAL --> S2;
}
```

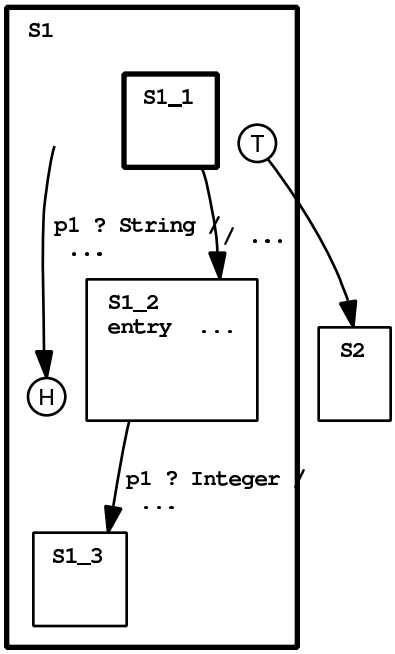

Example0025

```
package examples;

import org.echarts.ExternalPort;

public class Example0025Environment {

    public static final void main(String[] argv) {
        try {
            final ExternalPort p1 = new ExternalPort("p1");
            p1.input("forty two");
            p1.input(new Integer(42));
```

```
            new Example0025(p1).run();
        } catch (Exception e) { e.printStackTrace(); }
    }
}
```

The deep history pseudostate is graphically depicted by the letter 'H' surrounded by a circle. In this example, the machine's environment creates an external port `p1` and then enqueues two messages on the port's input queue: a string instance followed by an integer instance. Then the environment creates and runs the machine. The submachine defined in state `S1` fires its first transition from state `S1_1` to state `S1_2`, executing the entry action defined for `S1_2`, and then blocks in state `S1_2` waiting for an `Integer` instance to arrive on `p1`. Then the parent's first transition fires in response to the arrival of a `String` instance arriving on `p1` (see Section 6.4 for an explanation of the graphical depiction of this transition). The transition's target is the `DEEP_HISTORY` pseudostate of the `S1` submachine. Since the submachine was in state `S1_2` prior to the parent transition firing, then the submachine will return to state `S1_2`. Furthermore, since the `DEEP_HISTORY` reference does not constitute an explicit state reference, then the entry action defined for `S1_2` will not be executed again. Since the next message in `p1`'s input queue is an `Integer` instance, the submachine transition from `S1_2` to `S1_3` will fire. Finally, since `S1_3` is a terminal state, the parent's second transition will fire. For the record, here's the resulting output:

```
Hello World
Hola Mundo
forty two
42
```

The `DEEP_HISTORY` pseudostate will be discussed further when we discuss timed transitions in Section 3.16.

3.15.4 NEW

The `NEW` pseudostate was introduced when we discussed machine arrays in Section 3.14. When used to reference a machine array the effect is to create a new submachine element while leaving the state of any pre-existing machines undisturbed. In effect, a `NEW` pseudostate reference behaves exactly the same as a `DEEP_HISTORY` pseudostate reference for the pre-existing machines. Similarly, the `NEW` pseudostate behaves exactly the same as a `DEEP_HISTORY` pseudostate reference for and-machines or or-machines.

3.15.5 Summary

Here's a summary of how the pseudostates are interpreted by the ECharts runtime system.

Pseudostate	Source/Target	Implicit/Explicit
DEFAULT_INITIAL	target	explicit
TERMINAL	source	explicit
DEEP_HISTORY	target	implicit
NEW	target	implicit

3.16 Timed Transitions

Timed transitions are used for triggering timed events. A timed transition's timer is *activated* (starts ticking if it wasn't already ticking) when the state referenced by the transition's source becomes the current state. Here's a simple example:

```
package examples;

public machine Example0026 {
    initial state S1;
    state S2;
    transition S1 - delay(1000) / {
        System.out.println("duration: " + duration);
        System.out.println("activated: " + activationTime);
        System.out.println("expired: " + expiryTime);
    } -> S2;
}
```

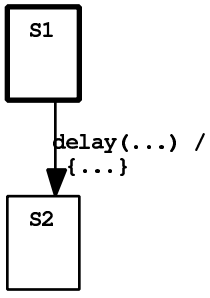

Example0026

All this machine does is pause in state `S1` for 1000 ms prior to transitioning to state `S2`. The transition's action prints out the values of three distinguished variables available during timed transition

actions: `duration`, `activationTime`, and `expiryTime`. `duration` is the duration in ms associated with the transition, `activationTime` is the timestamp in ms at which the transition's timer became activated, and `expiryTime` is the timestamp in ms at which the transition's timer expired. Note that the expiry time is not necessarily equal to the time that the transition fires since other transitions may fire in the interval between the transition timer's expiry and the transition firing. For further discussion concerning the relative priority of ports and transitions, see Sections 4.3 and 4.4.

Beyond being simply activated, a timed transition's timer will be *reactivated* (activated with its counter reset) if the state referenced by the transition's source is explicitly referenced by the previously firing transition's target. On the other hand, a timed transition's timer will remain activated or expired if referenced with a `DEEP_HISTORY` pseudostate (the `DEEP_HISTORY` pseudostate is discussed in more detail in Section 3.15.3). Here's an example of both of these concepts:

```
package examples;

public machine Example0027 {

    <* private long duration = 1000 *>
    <* private boolean fired = false *>

    initial state S1: {
        initial state S1_1: {
            initial state S1_1_1;
            state S1_1_2;
            transition S1_1_1 - delay(duration) /
                System.out.println("deep duration: " + duration)
            -> S1_1_2;
        }
        state S1_2;
        transition S1_1 - delay(duration) /
            System.out.println("shallow duration: " + duration)
        -> S1_2;
    }
    transition S1 - [ !fired ] / {
        duration = 1500;
        fired = true;
    } -> S1.S1_1.DEEP_HISTORY;
}
```

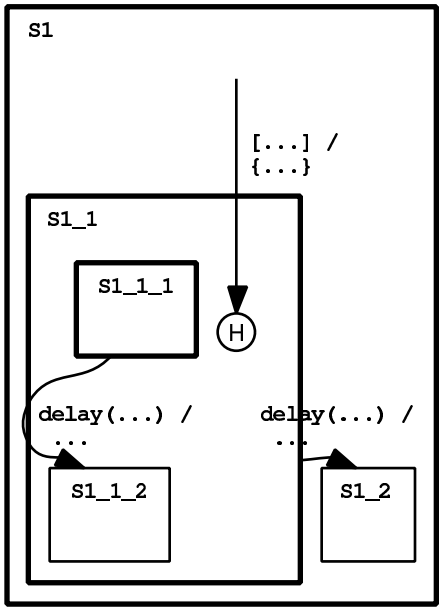

Example0027

In this example, the initial state of the machine activates the timers for the two timed transitions. Both timers are activated with the same duration value of 1000 ms. However, the first transition to fire is the messageless transition which changes the duration value. Because the transition's target state explicitly references state `S1.S1_1`, the timer for the (shallow) transition whose source state is `S1.S1_1` is reactivated with the new duration value (see Section 6.4 for an explanation of the graphical depiction of this transition). But the timer for the (deep) transition whose source state is `S1.S1_1.S1_1_1` remains activated with the original duration value because it is referenced with a `DEEP_HISTORY` pseudostate. Here's the machine's output:

```
deep duration: 1000
shallow duration: 1500
```

A timed transition whose duration is 0 ms will immediately be enabled. A timed transition whose duration has a negative value will never be enabled.

When a guard condition is specified for a timed transition, it is important that the guard's true and false conditions be handled by the transition itself using compound targets (see Section 3.7). This is because if the transition's timer expires and no target is defined for the transition then an exception will be raised. See Section 4.5.1 for a detailed explanation of this behavior.

The interested reader is referred to [1] for a formal description of the rules used to determine timed transition activation and reactivation.

3.17 Submachine Access

Submachine fields and methods can be accessed from transition guards, transition actions, and entry/exit actions. Every machine maintains variables for referencing its submachines. A submachine's variable name is the same as its state's name. Here's an example:

```
package examples;

public machine Example0028 {

    initial state S1: {
        <* final String field = "Hello" *>
        initial state S1_1: {
            <* final String field = "World" *>
            initial state S1_1_1;
        }
    }
    state S2;
    transition S1 - / {
        System.out.println(S1.field);
        System.out.println(S1.S1_1.field);
    } -> S2;
}
```

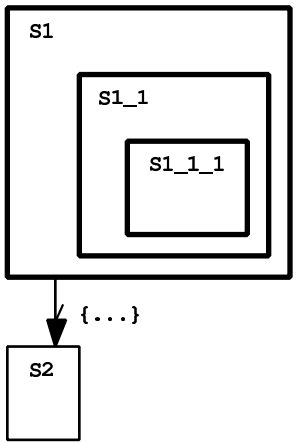

Example0028

In this example, the root machine accesses the fields named `field` in submachines defined for states `S1` and `S1.S1_1`. The machine's output is:

```
Hello
World
```

Here's another example, this time illustrating machine array element access (see Section 3.14).

```
package examples;

public machine Example0044 {
    <* private final int bound = 2 *>
    <* private int created = 0 *>
    <* private final String[] messages =
        new String[] { "Hello", "World!" } *>

    initial state S1[bound]: {
        <* private final String message =
            messages[getMachineIndex()] *>
        initial state S1_1;
    }
    transition S1 - [ created < bound ] /
        created++
    -> S1.NEW;
    <* private int index = -1 *>
    transition S1[?index].TERMINAL - [ created == bound ] /
        System.out.println("terminated: " +
            S1[index].message)
    -> S1;
}
```

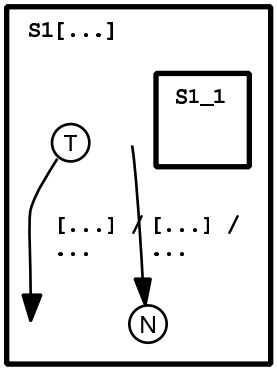

Example0044

Here's the machine's output:

```
terminated: Hello
terminated: World!
```

In this example, two machine array elements are initially created. The machine array elements consist only of a single state that plays the role of both the machine's initial state and its terminal state (see Section 3.15.2). Following their creation (see Section 6.4 for an explanation of the graphical depiction of this transition), the second transition is enabled fire for each element. The notation `S1[ ?index ].TERMINAL` has the effect of setting the `index` variable to the value of the machine array element index satisfying the transition's source state reference. The resulting `index` value is used to access the `message` field of the referenced machine array element in the transition's action.

We discuss precisely when submachine variable values are set and cleared in Section 4.6. Also, submachine variable access is constrained by machine and state access permissions. This topic is discussed in Section 4.8.

3.18 Partial State References

A machine transition's source and target state references need not specify a state at all. When no source state is referenced then the transition's source state is satisfied regardless of the current state of the machine. When no target state is referenced then the target state is treated as a reference to the machine's `DEEP_HISTORY` pseudostate (see Section 3.15.3). Here's an example:

```
package examples;
```

```
public concurrent machine Example0029 {

    <* final ExternalPort p1 *>

    public Example0029(ExternalPort p1) {
        this.p1 = p1;
    }
    <* private int messages = 0 *>
    <* private boolean end = false *>
    state S1: {
        initial state S1_1;
        state S1_2;
        transition S1_1 - p1 ? String [ !end ] / {
            messages++;
            System.out.println(message);
        } -> S1_1;
        transition S1_1 - [ end ] -> S1_2;
    }
    state S2: {
        initial state S2_1;
        state S2_2;
        transition S2_1 - p1 ? Integer [ !end ] / {
            messages++;
            System.out.println(message);
        } -> S2_1;
        transition S2_1 - [ end ] -> S2_2;
    }
    transition [] - [ !end && messages == 2 ] /
        end = true
    -> [];
}
```

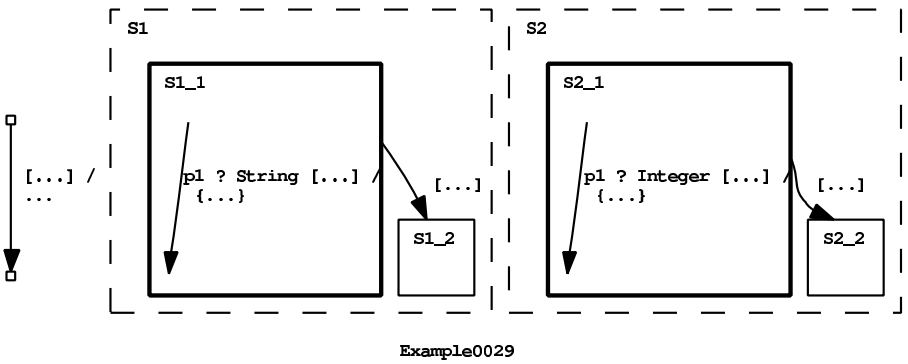

```
package examples;

import org.echarts.ExternalPort;

public class Example0029Environment {

    final public static void main(String[] argv) {
        try {
            final ExternalPort p1 = new ExternalPort("p1");
            p1.input("Hello");
            p1.input(new Integer(42));
            p1.input("World");
            new Example0029(p1).run();
        } catch (Exception e) { e.printStackTrace(); }
    }
}
```

In this example, the machine's environment enqueues three messages in external port `p1`'s input queue. This port is then passed to the machine as a constructor parameter. The machine maintains two concurrent submachines. One increments `messages` when a `String` instance arrives on `p1` and the other increments `messages` when an `Integer` instance arrives on `p1` (see Section 6.4 for an explanation of the graphical depiction of these transitions). When the number of messages to arrive on `p1` reaches two, the transition with no specified source or target state fires setting the flag that is referenced by the submachines, indicating that they should cease to wait for further messages on `p1`. As shown in the machine diagram, the graphical notation for an empty source or target state reference is a small empty box.

When we introduced join and fork transitions in Section 3.6 the examples showed all concurrent states of an and-machine being explicitly referenced. This is not necessary in general. Omitting a concurrent state in a transition's source is treated as a "don't care" reference. Omitting a concurrent state in the transition's target is treated as a `DEEP_HISTORY` reference. These points are illustrated in the following example:

```
package examples;

public concurrent machine Example0030 {
    state S1: {
        initial state S1_1;
        state S1_2;
```

```
    }
    state S2: {
        initial state S2_1;
        state S2_2;
    }
    transition S1.S1_1 --> S1.S1_2;
}
```

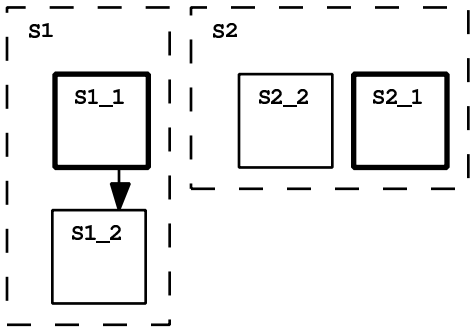

Example0030

In this example, the and-machine transition only explicitly references concurrent state `S1` (see Section 6.4 for an explanation of the graphical depiction of this transition). Since `S2` is not referenced in the transition's source state, then `S2` is guaranteed to satisfy the transition's source state. Furthermore since `S2` is not referenced in the transition's target state, then the reference to `S2` is treated as a `DEEP_HISTORY` pseudostate. When the transition executes, state `S1` changes from `S1_1` to `S1_2` and state `S2` remains in `S2_1`.

3.19 Host Language Interface

In Section 1 we introduced ECharts as a *hosted* language. This means that ECharts was never intended to be a complete programming language. Rather, ECharts relies on an underlying host language to support data operations and low-level control flow. ECharts strives to make the boundary between itself and its host language as seamless as possible while also trying to be independent from its host language. Experience gathered from earlier versions of ECharts has led us to adopt a boundary where basic language expressions are directly supported by the ECharts language, but more complex host control-flow constructs and field and method declarations are be contained within host language delimiters. Consistent with ECharts adoption of other Java concepts, ECharts has adopted Java-style expressions. These expressions are translated appropriately to the chosen target host lan-

guage. Escaped host constructs delimited by `<* *>`, are stripped of their delimiters and inserted directly into the translated code.

To this point, the manual has provided many examples of escaped host code constructs and unescaped ECharts expressions. It is easiest to describe which ECharts expressions are acceptable indirectly by enumerating those Java expressions that are *not* acceptable ECharts expressions and, therefore, must be wrapped in host language delimiters: `instanceof`, `if, then, else`, `try, catch, throw`, `switch, case`, `do, while`, `for`, `synchronized`, `return, continue, break, assert`, `?:` ternary operator, method declarations, variable declarations, and inner class declarations.

In the ECharts grammar a host language element wrapped in `<* *>` plays the role of a primary expression. As such it may be embedded in a larger ECharts expression. For example, a transition guard for a Java hosted machine might look like this:

```
[ <* isLongString(var1) ? var2 : var3*>.someMethod() == 42 ]
```

4 The Runtime Model

To this point we've examined ECharts' language features. What we've intentionally ignored so far are the fundamental semantics underlying the language. This section answers some of the questions that you may have been wondering about such as: When more than one transition can fire, which one is chosen? Can data be shared between peer machines? When are machines created and destroyed?

4.1 Machine Execution

As described in Sections 3.9 and 3.10, there are two ways to execute an ECharts machine: (1) a machine's execution blocks until a message arrives from the environment and (2) a machine method is invoked with a port and a message as arguments and, therefore, does not need to block. We describe the first mode of execution in more detail here. The non-blocking mode of execution is more complex so we omit its description for the sake of brevity.

Blocking execution of an ECharts machine proceeds as follows:

```
Fire a maximal sequence of maximum priority enabled
messageless transitions.
 while active message transitions exist in the machine's
current state:
 do
     while no messages enqueued for active ports:
     do
```

```
        Block execution until message arrives.
      end while
    Get a message from an maximum priority active port's
    input queue.
    Fire highest priority enabled message transition for the
    message and port.
    Fire maximal sequence of maximum enabled priority
    messageless transitions.
  end while
```

To understand the pseudocode above you should understand the distinction between active and enabled transitions. A transition is an *active transition* if its source state reference is satisfied by the machine's current state. A transition is an *enabled transition* if (1) it is an active transition and (2) the transition's guard conditions are satisfied. A firing transition must be an enabled transition. A port is an *active port* if it is specified by an active message transition. Each machine execution cycle, a message is removed from a currently active port. More discussion related to message dequeuing is found in Section 4.5. The concept of transition priority is discussed in Section 4.4. The concept of port priority is discussed in Section 4.3. Other issues related to execution blocking are discussed in Section 4.2.

4.2 Transition Scheduling

The essence of ECharts machine execution described in Section 4.1 is that the ECharts runtime attempts to fire as many transitions as possible in response to receiving a message from the environment. Only when there are no more transitions to execute is the next message from the environment considered. This approach to scheduling CPU resources is not uncommon for real-time systems. It is known as "run-until-blocked" or "run-to-quiescence" scheduling.

The sequence of one message transition followed by zero or more messageless transitions is referred to as a *transition sequence*. Note that the transitions in a transition sequence may belong to different submachines of a machine, including the submachines of concurrent states of an and-machine. However, there are constraints that dictate which transitions are considered for firing in a message sequence. This topic is addressed in Sections 4.3, 4.4 and 4.7.

A transition sequence is uninterruptible in the sense that once it is initiated for a machine no other messages from the environment will be acted upon until the sequence completes. Furthermore, the scheduling granularity is at the level of the transition. That is to say, a firing transition is guaranteed to run to completion without interruption.

There are a number of advantages afforded by this scheduling ap-

proach. The biggest one is that there is no need for explicit concurrency control to enforce data integrity constraints. For example, a machine's transition action that updates a data value need not take precautions to protect the data value from simultaneous update by other transitions belonging to the same machine. However, it is important to realize the extent of the protection afforded by scheduling applies only to actions associated with the machine itself. For example, if the aforementioned data value is accessed by other threads (other than the machine's thread) then the usual precautions should be taken, for example using Java's `synchronized` construct. Another advantage of this scheduling policy is that scheduling granularity is effectively controlled by programmer. The less time actions take to execute, the finer the scheduling granularity.

However, this scheduling approach also has potential disadvantages that go hand in hand with with its advantages: starvation and blocking. When a transition action executes for a long time, it effectively starves other transitions from executing. For this reason, it is important that the execution time of individual actions be as kept short as possible. Furthermore, when a machine's action makes a external blocking call no other transitions in that machine can fire until that blocking call returns. For example, if a state's entry action makes a blocking call to retrieve a value from an external database, then no machine transitions can fire until the blocking call returns. In some programming contexts extended external blocking does not not matter. However, when it does matter, then efforts should be taken to convert the (synchronous) blocking call to an (asynchronous) non-blocking call that signals its completion by sending a message from the environment to a machine port.

4.3 Port Priorities

In message receive operations, the class of ports implicitly associated with timed transitions have higher priority than external ports. In practice this means that if a timed transition has expired and is ready to fire, and a message is waiting in the input queue for an external port of a message transition that is ready to fire, then the ECharts runtime will fire the timed transition. In general, enabled timed transitions will always fire before enabled message transitions. Here's a simple example:

```
package examples;

public machine Example0031 {
    <* final ExternalPort p1 *>
```

```
    public Example0031(ExternalPort p1) {
        this.p1 = p1;
    }
    initial state S1;
    state S2;
    transition S1 - p1 ? String /
        System.out.println("message transition fired")
    -> S2;
    transition S1 - delay(0) /
        System.out.println("timed transition fired")
    -> S2;
}
```

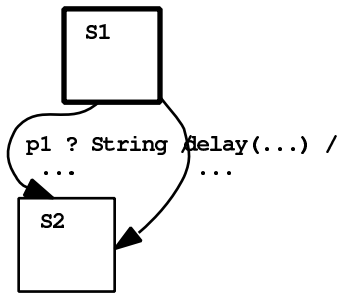

Example0031

```
package examples;

import org.echarts.ExternalPort;

public class Example0031Environment {

    public static final void main(String argv[]) {
        try {
            final ExternalPort p1 = new ExternalPort("p1");
            p1.input("Hello World!");
            new Example0031(p1).run();
        } catch (Exception e) { e.printStackTrace(); }
    }
}
```

In this example, the timed transition expires immediately. However, the message transition is also enabled because a message was

placed in the transition port's input queue by the machine's environment. Although both transitions are enabled, the timed transition is guaranteed to fire since its implicit port has higher priority than the message transition's port.

Support for internal ports will be provided in the next version of ECharts (see Section 9.3). Internal ports will have a priority between that of timed transition ports and external ports. Transition priority rules are discussed in Section 4.4.

4.4 Transition Priorities

An important advantage of ECharts relative to other Statecharts dialects is that ECharts provides programmers with a number of ways to control which transition will fire when more than one transition is enabled in a machine state. Section 4.3 discusses one such way based on relative port priorities. In addition, ECharts provides three rules to resolve relative transition priority. The three priority rules are applied in the order we present them here. If one rule does not resolve the priority between comparable transitions, then the next rule is applied.

4.4.1 Message Class Rule

The first priority rule applies only to message transitions. If more than one message transition is enabled for a message from a given port, then the transition specifying the most specific message class will be chosen to fire. Here's an example:

```
package examples;

public machine Example0032 {
    <* final ExternalPort p1 *>
    public Example0032(ExternalPort p1) {
        this.p1 = p1;
    }
    initial state S1: {
        initial state S1_1;
        state S1_2;
        transition S1_1 - p1 ? Object /
            System.out.println("Got Object: " + message)
        -> S1_2;
    }
    state S2;
    transition S1 - p1 ? String /
        System.out.println("Got String: " + message)
```

```
    -> S2;
}
```

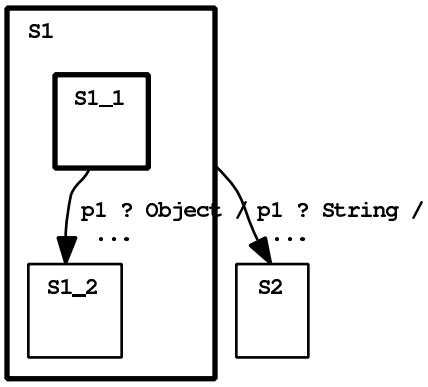

Example0032

```
package examples;

import org.echarts.ExternalPort;

public class Example0032Environment {

    public static final void main(String argv[]) {
        try {
            final ExternalPort p1 = new ExternalPort("p1");
            p1.input("Hello World!");
            new Example0032(p1).run();
        } catch (Exception e) { e.printStackTrace(); }
    }
}
```

Both transitions in the machine are enabled because the message input to the port is an instance of both the `Object` and `String` classes. However, because the `String` class specified in the root machine transition is a subclass of the `Object` class specified in state `S1`'s submachine, then the `String` transition fires.

This priority rule is most often used as a way to override default behavior. One transition, intended to handle default behavior, specifies a general message class, and others are defined with more specific message classes in order to override the default transition when necessary.

When an active port is shared amongst submachines of an and-machine (see Section 3.5) or submachine elements of a machine array (see Section 3.14), this priority rule is applied across all submachines guaranteeing that the transition chosen to fire will be the highest priority transition across all submachines. You should realize that this can lead to poor performance when there are many concurrent submachines since potentially all submachines must be inspected to locate the highest priority transition.

4.4.2 Source Coverage Rule

The second priority rule applies to all transitions for which the first rule fails to resolve priorities. The enabled transition with the highest priority is the transition with the most specific source state reference. This rule is a generalization of what is typically the only transition priority rule supported by other Statecharts dialects. Here's an example:

```
package examples;

public machine Example0033 {
    <* final ExternalPort p1 *>
    public Example0033(ExternalPort p1) {
        this.p1 = p1;
    }
    initial state S1: {
        initial state S1_1;
        state S1_2;
        transition S1_1 - p1 ? Object /
            System.out.println("Source S1_1: " + message)
        -> S1_2;
    }
    state S2;
    transition S1 - p1 ? String /
        System.out.println("Source S1: " + message)
    -> S2;
    transition S1.S1_1 - p1 ? String /
        System.out.println("Source S1.S1_1: " + message)
    -> S2;
}
```

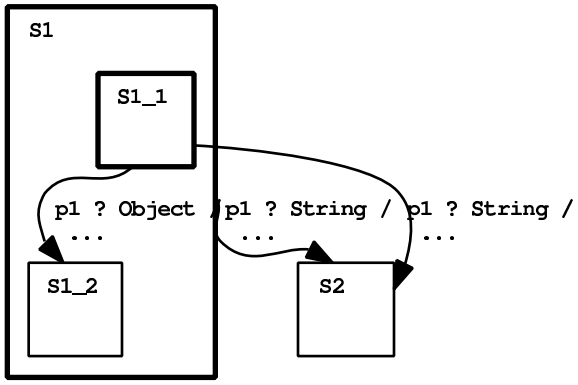

Example0033

```
package examples;

import org.echarts.ExternalPort;

public class Example0033Environment {

    public static final void main(String argv[]) {
        try {
            final ExternalPort p1 = new ExternalPort("p1");
            p1.input("Hello World!");
            new Example0033(p1).run();
        } catch (Exception e) { e.printStackTrace(); }
    }
}
```

And here's the machine's output:

```
Source S1.S1_1: Hello World!
```

In this case all message transitions in the machine are initially enabled. But the two transitions in the root machine have higher priority than the transition in state `S1`'s submachine because of the first priority rule discussed in Section 4.4.1. However, since the two root machine transitions have equal priority according to the first rule, then the second rule is applied. Since source state `S1.S1_1` referenced by the second transition is more specific than the source state `S1` referenced by the first transition, then the second transition has higher priority than the first transition under the second priority rule, so the second transition fires.

As mentioned earlier, the second rule applies not just to message transitions but also to messageless transitions. Here are two more examples.

```
package examples;

public machine Example0034 {
    initial state S1: {
        initial state S1_1: {
            initial state S1_1_1;
        }
        state S1_2;
        transition S1_1 - /
            System.out.println("Source S1_1")
        -> S1_2;
    }
    state S2: {
        initial state S2_1;
        state S2_2;
        transition S2_1 -/
            System.out.println("Source S2.S2_1")
        -> S2_2;
    }
    state S3;
    transition S1.S1_1.S1_1_1 - /
        System.out.println("Source S1.S1_1.S1_1_1")
    -> S2;
    transition S2 - /
        System.out.println("Source S2")
    -> S3
}
```

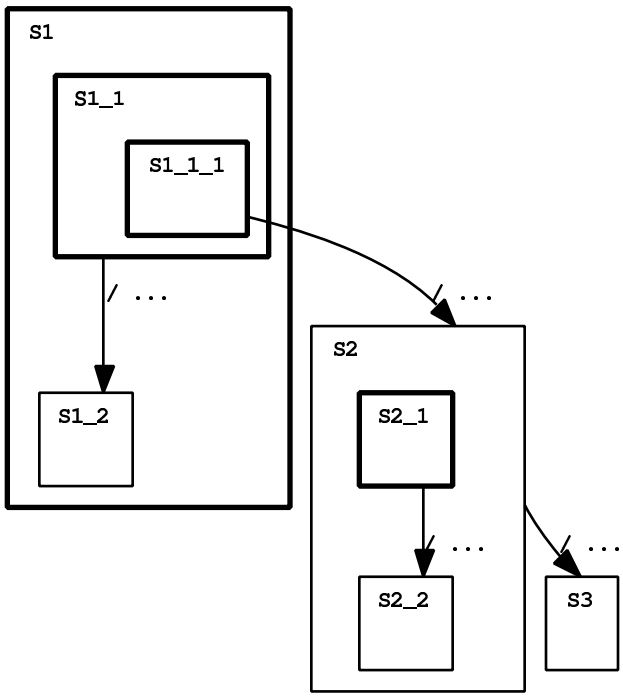

Example0034

Here is the machine's output:

```
Source S1.S1_1.S1_1_1
Source S2.S2_1
Source S2
```

In the initial state, the root machine transition referencing source state `S1.S1_1.S1_1_1` and the submachine transition referencing source state `S1_1` are enabled. However, because the transition in the root machine references a more specific source state than the submachine transition then the root machine transition fires. Two more transitions are enabled once the machine transitions to state `S2`: the root machine transition referencing source state `S2` and the submachine transition referencing source state `S2_1`. In this case, the submachine transition is guaranteed to fire first because its source state is more specific than the root machine transition's. Finally, the parent machine fires since it remains enabled after the submachine transition fires,

The second priority rule is useful for two purposes. One purpose is to override the default transition behavior defined for a machine source state. To do this the programmer defines additional transitions defined for particular source substates. This use of the rule is shown

in the last two examples. Another purpose is to allow submachine transitions to fire before more general ancestor machine transitions fire. This supports the view of a submachine invocation being akin to a procedure call in a traditional language. This use of the rule is shown in the last part of the second example above.

When the `TERMINAL` pseudostate (see Section 3.15.2) is referenced by a transition's source, then it is considered to be less specific than a reference to an actual machine state. The second priority rule also generalizes to join transitions (discussed in Section 3.6). In this case, each join transition branch reference must be less specific in order for the entire join to be considered less specific. The notion of source state *specificity* is formally captured by *term coverage* which we will not cover in this document in the interests of brevity. The curious reader is referred to [3] for more information.

4.4.3 Transition Depth Rule

This last rule is applied when the previous two rules fail to resolve the relative priority of two comparable transitions. Like the source coverage rule discussed in Section 4.4.2, this rule applies to all types of transitions: timed, messageless and message transitions. When the source states references of two transitions are identical, this rule gives highest priority to the transition defined in a higher-level machine (or, if you prefer, a shallower depth machine). Here's an example:

```
package examples;

public machine Example0035 {
    initial state S1: {
        initial state S1_1: {
            initial state S1_1_1;
        }
        state S1_2;
        transition S1_1.S1_1_1 - /
            System.out.println("Source S1_1.S1_1_1")
        -> S1_2;
    }
    transition S1.S1_1.S1_1_1 - /
        System.out.println("Source S1.S1_1.S1_1_1")
    -> S1.S1_2;
}
```

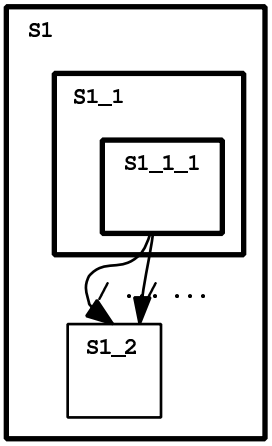

Example0035

Here's the machine's output:

```
Source S1.S1_1.S1_1_1
```

In this example both transitions refer to the same source state, namely state `S1.S1_1.S1_1_1`. In this case, the application of the transition depth rule guarantees that the root machine's transition fires (see Section 6.4 for an explanation of the graphical depiction of this transition), since the root machine depth is shallower (depth 0) than the depth of the submachine defined for state `S1` (depth 1).

The motivation for using this rule is to override submachine behavior. If a programmer would like to alter the behavior of a particular submachine's transition, then a transition in an ancestor machine can be defined to reference the same source state and, optionally, the same target states as the submachine transition. This priority rule ensures that the ancestor transition will fire instead of the submachine's transition.

4.5 Message Dequeuing

In Sections 3.9 and 3.12 devoted to receiving and sending messages, we discussed how an external port maintains an input queue into which messages from a machine's environment are enqueued. In Section 4.3 we discussed how dequeuing messages from timed transition ports has higher priority than dequeuing messages from other message transition ports. In this section we discuss two other issues related to dequeuing messages. The ECharts approach to message dequeuing attempts to balance two conflicting desires: (1) to never unintentionally lose

messages from the environment and (2) to not overly burden the programmer. These two issues are addressed, in turn, in the next two sections.

4.5.1 Explicit Message Consumption

The ECharts runtime model guarantees that no messages from the environment are unintentionally lost. The justification for this is simply that in most reactive programming domains, it is unacceptable to unknowingly lose a message. This aspect of ECharts distinguishes it from other Statecharts dialects that typically allow a message to be lost if it is not explicitly handled by a machine.

Here's how ECharts ensures that messages are not lost: if p is an active port (see Section 4.1), then the ECharts runtime may dequeue a message from port p. Furthermore, if a message is dequeued from port p, then the machine *must* fire some message transition for that message, otherwise an exception is raised. Here's an example:

```
package examples;

public machine Example0036 {
    <* final private ExternalPort p1 *>
    public Example0036(ExternalPort p1) {
        this.p1 = p1;
    }
    initial state S1;
    state S2;
    transition S1 - p1 ? String -> S2;
}
```

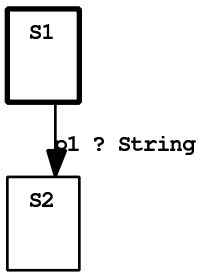

Example0036

```
package examples;
```

```
import org.echarts.ExternalPort;

public class Example0036Environment {
    final static public void main(String[] argv) {
        try {
            final ExternalPort p1 = new ExternalPort("p1");
            p1.input(new Integer(42));
            new Example0036(p1).run();
        } catch (Exception e) { System.out.println(e); }
    }
}
```

Here's the machine's output:

```
org.echarts.MachineException: No transition defined for
message class java.lang.Integer with string representation 42
on port PORT_NAME=p1, PORT_TYPE=EXTERNAL_PORT,
PORT_ID=192.168.1.102:5420:109c366d015:-8000
from state :examples.Example0036 S1 (S2)
```

In this example, the machine's environment enqueues an `Integer` instance on port `p1`'s input queue. However, the machine only defines a message transition for a `String` instance on `p1`, so the ECharts runtime raises an exception that is caught by the environment. If there had been an additional message transition for `p1` that specified the `Integer` class or the `Object` (super) class then no exception would have been raised.

4.5.2 Implicit Message Deferral

Having seen how unhandled messages are treated in the previous section, you may be asking yourself if it is necessary to add message transitions to every machine state for every message class anticipated on every port. The answer to this question is an emphatic: No. This would be an unreasonable burden on the programmer. Furthermore, it would obfuscate a machine's logic. Instead, ECharts provides programmers with a facility to implicitly defer dequeuing messages from a port until the programmer explicitly declares a machine's readiness to accept a message from the port. The way this is accomplished is very simple: if a port p is not active (see Section 4.1), then the ECharts runtime will not dequeue a message from p. Only when the machine arrives in a state where p is active may a message be dequeued from p. If a message is dequeued from p then the machine is obliged to fire a transition for that message. Here's an example:

```
package examples;

public machine Example0037 {
    <* final ExternalPort p1 *>
    public Example0037(ExternalPort p1) {
        this.p1 = p1;
    }
    initial state S1;
    state S2;
    state S3;
    transition S1 --> S2;
    transition S2 - p1 ? String /
        System.out.println(message)
    -> S3;
}
```

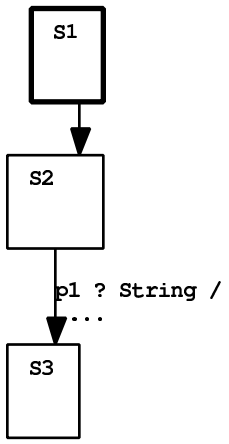

Example0037

```
package examples;

import org.echarts.ExternalPort;

public class Example0037Environment {

    final static public void main(String[] argv) {
        try {
            final ExternalPort p1 = new ExternalPort("p1");
            p1.input("Hello World");
            new Example0037(p1).run();
```

```
        } catch (Exception e) { e.printStackTrace(); }
    }
}
```

The example shows the machine environment enqueuing a message on port `p1`. However, the initial state of the machine does not include transitions specifying `p1` so the message remains enqueued on the port until the machine reaches state `S2` where there is a transition specifying `p1`. Since this transition specifies a `String` message then it fires and the machine completes its execution.

4.6 Machine Creation and Destruction

In the interests of conserving memory, ECharts machines are created only when they are needed and they are destroyed when they are no longer needed. This section explains in more detail what is meant by "needed" and "no longer needed."

4.6.1 Machine Creation

The root machine is created by the machine's environment, as we have seen in many of the previous examples. However, submachines can be created in one of three ways: (1) If a submachine does not already exist then it will be created when it becomes part of a machine's current state. (2) If a submachine already exists, then it will be *re-created* if its parent state is explicitly referenced as a target of a firing transition, and no sub-states of the parent state are referenced. (3) Finally, a machine array element can be created using the `NEW` pseudostate as discussed in Section 3.15.4. Here is an example of (1) and (2):

```
package examples;

public machine Example0038 {

    initial state S1: {
        <* private int field = 42 *>

        initial state S1_1;
        state S1_2;
        transition S1_1 - /
            System.out.println(field)
        -> S1_2;
    }
    <* private boolean fired = false *>
```

```
    transition S1 - [ !fired ] / {
        fired = true;
        S1.field = 0;
    } -> S1;
}
```

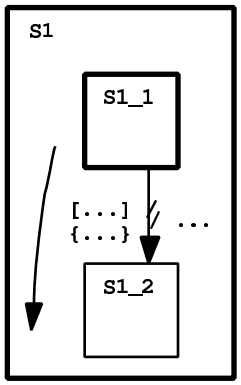

Example0038

Here's the machine's output:

```
42
42
```

In `Example0038`, the submachine defined for initial state `S1` is automatically created when the root machine is initially created. This is because `S1` is the initial current state of the machine. The submachine transition fires printing out the initial value of `field`: 42. Then the parent transition fires (see Section 6.4 for an explanation of the graphical depiction of this transition). Since the transition target explicitly references `S1` and no substates of `S1`, then `S1`'s submachine is re-created. The evidence of this is that, although the parent transition changes the value of the submachine's field to 0, the value of submachine's field printed when the submachine transition fires again is its initial value: 42.

In contrast, consider the following example:

```
package examples;

public machine Example0039 {

    initial state S1: {
```

```
        <* private int field = 42 *>

        initial state S1_1;
        state S1_2;
        transition S1_1 - /
            System.out.println(field)
        -> S1_2;
    }
    <* private boolean fired = false *>

    transition S1 - [ !fired ] / {
        fired = true;
        S1.field = 0;
    } -> S1.DEFAULT_INITIAL;
}
```

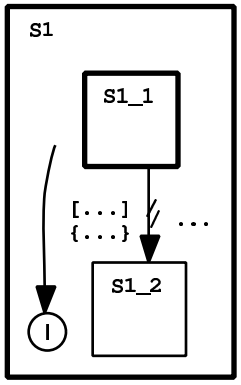

Example0039

This example is identical to the previous example except that the root transition target references `S1`'s `DEFAULT_INITIAL` pseudostate, not simply `S1` as before (see Section 6.4 for an explanation of the graphical depiction of this transition). In this case, `S1`'s submachine is not re-created since the transition references a sub(pseudo)state of `S1`. This is evident from the program's output:

```
42
0
```

The rules for machine creation ensure that a machine always exists when you expect it to. The need for machine re-creation is not common in practice, nevertheless it is a useful programming tool to have available.

The previous examples showed how the submachines of an or-machine's current state are guaranteed to exist. This generalizes to and-machines in the expected way. Since every state of an and-machine is part of the machine's current state, then all submachines of an and-machine's states are guaranteed to exist.

4.6.2 Machine Destruction

A machine can only be destroyed once it enters a terminal state (the concept of a terminal state is defined in Section 3.15.2). However, a machine is not destroyed upon entering a terminal state; it can only be destroyed by an ancestor machine. In particular, the source state of a transition in an ancestor machine must explicitly reference the descendant machine's terminal state in order for the descendant machine to be destroyed. Furthermore, the descendant machine isn't destroyed until the transition's actions have completed, thereby providing an opportunity to access the descendant machine's state immediately prior to its destruction. Here's an example:

```
package examples;

public machine Example0040 {
    initial state S1: {
        <* private int field = 42 *>
        initial state S1_1;
        state S1_2;
        transition S1_1 --> S1_2;
    }
    state S2;
    <* private boolean fired = false *>
    transition S1.S1_2 - [ !fired ] / {
        System.out.println("S1.field: " + S1.field);
        S1.field = 0;
        fired = true;
    } -> S1.DEEP_HISTORY;
    transition S1.TERMINAL - [ fired ] /
        System.out.println("S1.field: " + S1.field)
    -> S2;
}
```

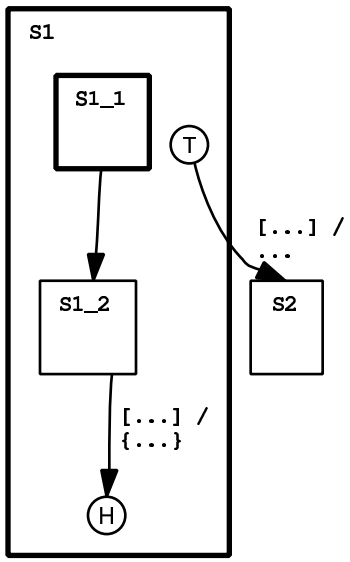

Example0040

Here's the machine's output:

```
S1.field: 42
S1.field: 42
```

In this example, the submachine in state `S1` transitions to a terminal state. Then the first root machine transition fires (see Section 6.4 for an explanation of the graphical depiction of this transition), printing the initial value of the submachine's `field` variable, changing its value to 0, and then looping back to the `DEEP_HISTORY` pseudostate of the submachine. Since the submachine was in an terminal state, and since the root machine transition explicitly referenced the terminal state, then the submachine will have been destroyed when the transition fired. This means that the changed value of the submachine's `field` variable will be lost when the submachine is destroyed. The transition's target is state `S1`, so `S1` becomes the machine's current state again. As described in the previous section, Section 4.6.1, this means that a new instance of the submachine in `S1` must be created. Since the newly created submachine has no prior state, the transition target `DEEP_HISTORY` pseudostate reference defaults to the submachine's initial state `S1_1`. Once again, the submachine's transition fires, moving the submachine to a terminal state. Now the second transition in the root machine fires, printing out the current value of the submachine's `field` variable. Since the variable is defined in a newly created submachine instance then the initial field value, 42, is printed. You should

also note that since the transition source references the submachine with the `TERMINAL` pseudostate, then the submachine will, once again, be destroyed. This time, however, the new current state of the machine becomes `S2`, so the submachine in state `S1` is not re-created.

The same rule for machine destruction applies to elements of a machine array (see Section 3.14). In the context of a machine array, it is important for the programmer to remember to include transitions to destroy machine array elements that have arrived in a terminal state. Neglecting to do this will eventually result in the machine array reaching its maximum capacity leaving no opportunity for new elements to be added. Furthermore, this situation may also over-utilize memory resources.

The destruction of a machine can be prevented by annotating the machine's parent state with the `nonterminal` state modifier. As an example, see `Example0022` in Section 3.14. This modifier informs the ECharts interpreter to treat the parent state as a non-terminal state even though it might be a terminal state.

4.7 Shared Data

As described in Section 3.17, sharing data amongst ancestor and descendant machines is supported by ECharts. In the following sections we discuss constraints on data sharing.

4.7.1 `shared` Machine Modifier

Data sharing amongst peer concurrent submachines is discouraged for two reasons: (1) it is akin to sharing global data which reduces machine modularity, re-usability and maintainability, and (2) supporting it imposes a CPU performance penalty. Nevertheless, until internal ports (see Section 9.3) and past and future state references (see Section 9.4) are supported by ECharts, we recognize that it may occasionally be necessary to share data amongst peer submachines. Support for this is provided by the `shared` machine modifier. This modifier is applicable only for machines with concurrent states: and-machines (see Section 3.5) and machine arrays (see Section 3.14). Including this modifier modifies how the ECharts runtime searches for enabled messageless transitions to fire.

To understand how the `shared` modifier modifies the behavior of the ECharts runtime system, you should first understand its default behavior for choosing transitions to fire in a transition sequence (see Section 4.2). When a transition fires in a submachine of a concurrent state, the ECharts runtime system normally searches for a subsequent messageless transition to fire only in the submachine, its ancestor machines or its descendant machines. It does not search for messageless

transitions in submachines of peer concurrent states. Therefore, if a transition in a concurrent state's submachine were to modify a variable whose value is evaluated by a transition guard in a peer state's submachine, then the peer transition would not normally be considered for firing as part of the transition sequence. Adding the `shared` machine modifier modifies the runtime behavior so that all peer state submachines are checked for enabled transitions. Obviously, this increases CPU utilization relative to the default unshared behavior.

Here are two examples with machine arrays:

```
package examples;

public machine Example0041 {

    <* final private int bound = 2 *>
    <* private int created = 0 *>
    initial state S1[bound]: {
        initial state S1_1;
        state S1_2;
        transition S1_1 - [ created == bound ] /
            System.out.println("S1 index: " + getMachineIndex())
        -> S1_2;
    }
    transition S1 - [ created < bound ] /
        created++
    -> S1.NEW;
}
```

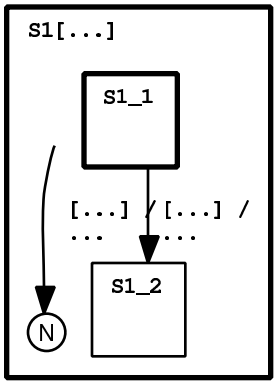

Example0041

The output from this example is:

```
S1 index: 1
```

In example `Example0041` above, the `shared` modifier is not used in the declaration of the machine array defined in state `S1`. When the first submachine is created (see Section 6.4 for an explanation of the graphical depiction of this transition) its transition does not fire because its guard is not satisfied. When the second submachine is created, however, its transition does fire because its guard is satisfied. This is indicated by the machine's output. At this point, the transition's guard in the first submachine is satisfied but the ECharts runtime does not choose the transition to fire because the previous transition to fire was in a submachine of a peer concurrent state. Now consider an example using the `shared` modifier.

```
package examples;

public machine Example0042 {

    <* final private int bound = 2 *>
    <* private int created = 0 *>
    initial state S1[bound]: shared {
        initial state S1_1;
        state S1_2;
        transition S1_1 - [ created == bound ] /
            System.out.println("S1 index: " + getMachineIndex())
        -> S1_2;
    }
    transition S1 - [ created < bound ] /
        created++
    -> S1.NEW;
}
```

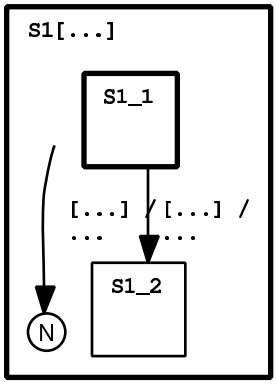

Example0042

The output from this example is:

```
S1 index: 1
S1 index: 0
```

The only difference between `Example0042` above and the previous example is the use of the `shared` machine modifier for the machine array. The result of using this modifier is that the enabled transition in the first submachine fires.

4.7.2 Port Variables

A message transition specifies a port variable. The ECharts runtime determines the value of the port variable when a transition becomes active. An important constraint on message transitions is that only the machine in which the transition is defined may change the value of the transition's port. If one machine changes a port variable's value in another machine (using submachine access discussed in Section 3.17 or using shared variables discussed in Section 4.7.1) then there is no guarantee that the correct port value will be utilized by a message transition in the other machine. With the introduction of internal ports in the near future (see Section 9.3) it will be straightforward for a machine to effect the change of another machine's port variable value. However, currently if one machine wishes to change the value of another machine's port variable, then it should use a shared variable to signal the other machine that its port variable has been changed.

4.8 Machine and State Access Modifiers

Access permissions can be specified for ECharts machines, machine constructors and machine states. The permissions model used by

ECharts is the same as that used by Java. There are four permission classes: `private` (currently unsupported for machines), `public`, `protected` (currently unsupported), and package (the default permission class if no permission class is explicitly specified). Machine pseudostates are assigned `public` access.

The permissions for machines, machine constructors and machine states are interpreted the same was as for Java class definitions, instance constructors, and instance methods, respectively.

Machine permissions and state permissions dictate submachine access permissions (see Section 3.17) and submachine state reference permissions in multi-level transitions (see Section 3.3). Machine permissions and machine constructor permissions dictate permissions for specifying and creating external submachines (see Section 3.1.2).

5 The Machine Runtime

An ECharts runtime executes ECharts machines (see Section 2.1). There are a number of configuration options and classes shared by ECharts runtimes to support runtime tasks such as monitoring and debugging. These are examined in the following sections.

5.1 Initialization

All ECharts runtime options can be configured when the runtime is initialized. There are two techniques for configuring options when the runtime is initialized: (1) using a properties file and (2) using command-line options. Some of these options can also be configured programmatically after initialization by invoking methods defined by the ECharts runtime API.

If option configuration techniques are combined then the following precedence order is obeyed: method invocation overrides command-line options, and command-line options override properties file settings.

For the *javamachine* runtime, a properties file takes the usual Java properties file format where each line in the file is of the form:

```
option: value
```

Command-line options for the *javamachine* runtime take the usual format for defining Java virtual machine system properties with the `java` command, namely:

```
-Doption=value
```

We discuss the various options in detail in the following sections.

5.2 Properties File

If the ECharts runtime finds a file named *echarts.properties* in the runtime's working directory then it will use the properties definitions found in that file. The runtime command line option `org.echarts.properties.dir` may also be used to specify a path to a directory other than the working directory in which to find the *echarts.properties* file.

5.3 Startup Messages

When an ECharts runtime is initialized, the initialization options can be printed to standard output by setting the `org.echarts.system.startupMessages` option. This option is disabled (`false`) by default but can be enabled by setting the option to `true`. This option can only be configured from a properties file or from the command-line.

5.4 Transition Timer Manager

An ECharts runtime uses a default transition timer manager to support timed transitions. However, for some application domains it is appropriate to use a domain-specific transition timer manager. For example, the ECharts interface for SIP servlet containers defines its own transition timer manager. The `org.echarts.system.transitionTimerManager.class` property specifies a fully qualified class name for the ECharts runtime to use as its transition timer manager. The specified class must implement the `org.echarts.TransitionTimerManager` interface. The default value for this property is `org.echarts.DefaultTransitionTimerManager`. The runtime loads the specified class and creates an instance of it for use by the runtime. This option can only be configured from a properties file or from the command-line.

5.5 Monitoring and Logging

An ECharts runtime includes a simple, configurable, extensible monitoring subsystem. This subsystem receives events generated by the runtime, optionally filters out particular events, and then formats and logs the events. The runtime uses this subsystem to support debugging, however, it can be used for many other purposes such as monitoring application-specific events. It can also easily be integrated with an application's existing monitoring subsystem.

To log an event, an ECharts application must send it to a monitor (see Section 5.5.1). For example, the built-in ECharts trace debugger (see Section 5.6) sends a series of events to ECharts machine monitors

during the course of a debugging session. Events are submitted to a monitor via the monitor's `putEvent()` method (see Section 5.5.3). This method checks if the event passes through the monitor's event filter (see Section 5.5.4) and, if it does, then it logs the event using the monitor's formatter (see Section 5.5.5).

5.5.1 Monitors

A runtime's monitoring subsystem is disabled by default. When monitoring is disabled, no events are sent to the monitoring subsystem. To enable it, set the `org.echarts.system.monitoring` property to `true`, or invoke the `setSystemMonitoring()` method. In addition to enabling monitoring for the runtime, monitoring must be enabled for the individual machines from which you want to receive events. The simplest way to do this is to set the `org.echarts.machine.monitoring` property to `true` or invoke the `setDefaultMachineMonitoring()` method. This value is set to `false` by default. The value of this property is used as the default monitoring setting for newly created machines. This default value can be overridden for a particular machine by calling the `setMachineMonitoring()` method. The default behavior is that a child machine inherits the monitoring setting of its parent, so if you wish to use different monitoring settings in a parent and its child, you must use the `setMachineMonitoring()` method.

Each machine managed by an ECharts runtime has a monitor. Machines may share a monitor, or they may use different monitors. The `org.echarts.machine.monitor.class` property can be used to specify a fully qualified classname of a monitor. The specified class must implement the `org.echarts.monitor.MachineMonitor` interface. The runtime will load the class and create an instance of it to use as the default monitor for newly created machines. The default value for this property is `org.echarts.monitor.PrintStreamMonitor` (discussed in Sections 5.5.5 and 5.6). This default value can be overridden for a particular machine by calling the `setMachineMonitor()` method. The default behavior is that a child machine inherits the monitor of its parent, so if you wish to use different monitors between parent and child, you must use the `setMachineMonitor()` method.

For a discussion of monitor formatters and logging, see Section 5.5.5.

5.5.2 Events

All monitor events are instances of the ECharts `MachineMonitorEvent` class. When an event instance is created it is automatically timestamped. An event also possesses a properties list of string key/value pairs. When an event is logged, as described in 5.5.5, then the event's timestamp and properties list are printed.

Apart from the pre-defined trace debugger event classes, there is one additional pre-defined event class, the `InfoEvent` class. This event class can be used for a number of purposes by the programmer, for example, logging custom debug messages or exception stack traces. However, it is also possible to create custom event classes by simply subclassing the root `MachineMonitorEvent` class, or the `InfoEvent` class itself.

5.5.3 putEvent()

To submit an event to a monitor, the monitor's `putEvent()` method must be invoked. A `putEvent()` method is also defined for each machine instance. The programmer may choose which of these `putEvent()` methods to invoke. Invoking a machine's `putEvent()` method causes the input event to be augmented with four properties prior to calling the machine monitor's `putEvent()` method. These additional properties are intended to associate the event with the machine producing the event.

The `MACHINE_ID` and `ROOT_MACHINE_ID` properties specify the unique ID assigned to the machine instance producing the event and its root machine, respectively. The `MACHINE_STATE_PATH` property specifies the state path from the root machine to the machine producing the event. The `SEQUENCE_ID` property specifies the unique ID assigned to the transition sequence that the event is a member of (see Section 4.2 for a description of transition sequences). A sequence ID is actually the unique ID assigned to the port associated with the message transition that initiates the sequence.

If it is not desirable to augment events with these properties, then a machine monitor's `putEvent()` method should be invoked directly rather than indirectly invoking it via the machine's `putEvent()` method.

5.5.4 Event Filters

When a machine submits an event to its monitor via its `putEvent()` method, the monitor first passes the event through its event filter. If the event passes through the filter the event is logged, otherwise it is discarded. Defining an event filter for a monitor provides flexible, fine-grained control of a monitor's logging levels.

An event filter is an instance of the `MachineMonitorEventFilter` class. A filter string is specified in the class constructor. A filter string specifies passing/discarding an event based on its (sub-)class or on its property values. A null or empty filter string passes all events. For example, here is an event string that passes `InfoEvent` instances whose `MESSAGE` property value is `Hello`, and discards `InfoEvent` instances whose `MESSAGE` property is `World`:

```
+org.echarts.monitor.InfoEvent?MESSAGE=(Hello) |
-org.echarts.monitor.InfoEvent?MESSAGE=(World)
```

The filter string syntax supports more complex filters including the specification of regular expressions for property values and conjuncts of event classes. A complete description can be found in the `MachineMonitorEventFilter` class documentation.

The default event filter for the ECharts runtime's default monitor (see Section 5.5.1) can be set during initialization with the `org.echarts.machine.monitor.filterString` property, or at runtime with the `setDefaultMachineMonitorFilter()` method. The default value for this property is the empty string, which passes all events.

5.5.5 Formatters and Logging

Event logging refers to what a monitor does with an event once it passes through the monitor's event filter. Typically, the event is logged to a file. Event formatting refers to how an event is portrayed in the event log. The logging task is performed by the monitor itself while the formatting task is performed by a monitor's formatter. An ECharts runtime comes with three general purpose monitor classes: (1) a `NullMonitor` that effectively discards all events sent to it, (2) a `PrintStreamMonitor` that logs formatted events to a specified print stream, for example, a file or standard output, and (3) a `RawMonitor` that logs unformatted events in their raw, binary format to a specified file. An ECharts runtime also comes with two formatter classes: (1) a `PrettyPrintFormatter` which formats events as easy-to-read strings, and (2) a `TextFormatter` which simply stringifies events in a machine-readable format. We will describe the output of the `PrettyPrintFormatter` in more detail in Section 5.6 devoted to the trace debugger.

A `PrintStreamMonitor` using a `PrettyPrintMonitor` is the default value for the ECharts runtime default monitor. Section 5.5.1 describes how the default monitor value can be set.

Programmers are free to extend any of the monitor or formatter classes for their own purposes in order to, for example, integrate ECharts runtime monitoring with an existing logging subsystem, or to customize runtime monitoring to suit their application needs.

5.6 Debugging

An ECharts runtime supports trace debugging. When enabled, trace debugging generates a sequence of monitor events reflecting machine execution steps, namely, transition events and machine creation or destruction events.

5.6.1 Output Format

The actual format of the events depends on the event formatter utilized by a machine's monitor (see Section 5.5.5). In the examples shown in this section we will use the `PrettyPrintFormatter`. Here's the complete trace debugger output produced when running the "Hello World" example from the beginning of this document (Section 2.1).

```
time: 2006.03.31 23:12:22:141 EST (1143864742141)
event: MessagelessTransitionEvent
root: a42792:10a53a6d0e6:-8000
machine: a42792:10a53a6d0e6:-8000
state path: :(Example0001)
sequence: initializing
transition: [S1] --> [S2]
local state:
    ...
    S2
    ...
```

The `time` and `event` fields specify the time the event was created and the event's class, respectively. Transition events are not created until immediately *after* the associated transition has fired. Since trace debug events are submitted to a machine's `putEvent()` method, they include the four properties discussed in Section 5.5.3. The `PrettyPrintFormatter` formats these properties as the `root`, `machine`, `state path`, and `sequence` fields.

In this example, the `root` and `machine` fields specify the same machine since the root machine fired the transition. Since there are no submachines in this example, the `state path` simply specifies the root machine. Since the firing transition belongs to the initial sequence of messageless transitions that fire when a root machine is initialized by the runtime the ID is shown as `initializing`.

The next two fields are common to message transition and messageless transition events. The `transition` field specifies the source and target state references of the messageless transition that fired. Had the transition been a message transition, then the transition's specified message class and a string representation of the message satisfying the transition are also displayed. In the case that a transition defines a compound target (see Section 3.7), then only the chosen target state reference is displayed. Finally, the `local state` field specifies the updated state of the machine in which the transition fired. Only states explicitly referenced by the transition's target state are displayed. In the event that the affected states contain submachines, then the submachine states are also shown.

Now here's a more complex example. First we show the machine and its environment. Then we show the trace debugger output resulting from running the machine's environment.

```
package examples;

public machine Example0046 {

    <* final private ExternalPort p1 *>

    public Example0046(ExternalPort p1) {
        this.p1 = p1;
    }
    initial state S1: {
        initial state S1_1;
        state S1_2;
        transition S1_1 - p1 ? String /
            putEvent(new InfoEvent(message))
        -> S1_2;
    }
    state S2;
    transition S1.TERMINAL --> S2;
}
```

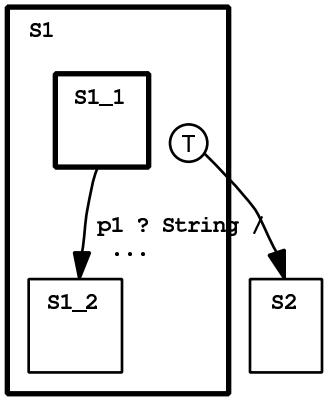

Example0046

```
package examples;

import org.echarts.ExternalPort;
```

```
public class Example0046Environment {

    public static final void main(String[] argv) {
        try {
            final ExternalPort p1 = new ExternalPort("p1");
            p1.input("Hello World!");
            new Example0046(p1).run();
        } catch (Exception e) { e.printStackTrace(); }
    }
}
```

```
time: 2006.03.31 23:48:38:154 EST (1143866918154)
event: MachineLifecycleEvent
root: ac2f9c:10a53c803f8:-7fff
machine: ac2f9c:10a53c803f8:-7fff
state path: :(Example0046)
sequence: initializing
created: S1:(inner machine)
submachine: ac2f9c:10a53c803f8:-7ffe
############################################################
time: 2006.03.31 23:48:38:221 EST (1143866918221)
event: InfoEvent
root: ac2f9c:10a53c803f8:-7fff
machine: ac2f9c:10a53c803f8:-7ffe
state path: :(Example0046).S1
sequence: ac2f9c:10a53c803f8:-8000
message: Hello World!
############################################################
time: 2006.03.31 23:48:38:224 EST (1143866918224)
event: MessageTransitionEvent
root: ac2f9c:10a53c803f8:-7fff
machine: ac2f9c:10a53c803f8:-7ffe
state path: :(Example0046).S1
sequence: ac2f9c:10a53c803f8:-8000
transition: [S1_1] - p1 ? String (Hello World!) -> [S1_2]
local state:
    ...
    S1_2
    ...

############################################################
time: 2006.03.31 23:48:38:249 EST (1143866918249)
event: MachineLifecycleEvent
root: ac2f9c:10a53c803f8:-7fff
```

```
machine: ac2f9c:10a53c803f8:-7fff
state path: :(Example0046)
sequence: ac2f9c:10a53c803f8:-8000
destroyed: S1:(inner machine)
submachine: ac2f9c:10a53c803f8:-7ffe
############################################################
time: 2006.03.31 23:48:38:281 EST (1143866918281)
event: MessagelessTransitionEvent
root: ac2f9c:10a53c803f8:-7fff
machine: ac2f9c:10a53c803f8:-7fff
state path: :(Example0046)
sequence: ac2f9c:10a53c803f8:-8000
transition: [S1.TERMINAL] --> [S2]
local state:
    ...
    S2
    ...

############################################################
```

The trace debugger output from this example shows two `MachineLifecyleEvents`. These events are associated with submachine creation and destruction. The `created` and `destroyed` fields indicate the parent state of the submachine being created or destroyed, respectively. It also shows the submachine's class if the submachine is an external machine, otherwise it shows `inner machine` as illustrated above. The `submachine` field specifies the submachine instance's unique ID. The trace also shows how an `InfoEvent` is displayed (see Section 5.5.2). Using `InfoEvent`'s is preferable to using a host language print statement because the event is logged and correlated in time with other trace debugger events. Furthermore, if the event was submitted to a machine's `putEvent()` method as described in Section 5.5.3 then the event will be automatically augmented with a number of additional properties.

The examples above show the default output format produced by the `PrettyPrintFormatter`. The `PrettyPrintFormatter` also has methods to control the display of debugger events including how machine class names and machine states are portrayed. See the `PrettyPrintFormatter` class API for details.

5.6.2 Message Properties

A message transition event includes a string representation of the message that triggered the transition by default. For the *javamachine* runtime, the string is the result of calling the message object's `toString()`

method. However, it is possible to include more structured information about the message in the form of a properties list of string key/value pairs. Not only is the properties list maintained in a structured format in a monitor log, the individual keys and values can also be referenced in a event filter in a straightforward fashion (see Section 5.5.4). For the *javamachine* runtime, a message class should implement the `org.echarts.monitor.Message` interface to associate a properties list with a message. When a message transition event is created for a message implementing this interface, the message properties are added to the event's properties. For an example, see the `org.echarts.TransitionTimeoutMessage` class. An instance of this class is sent by the ECharts runtime to a timer port when a timed transition expires. The message specifies three properties: `DURATION`, `ACTIVATION_TIME`, and `EXPIRY_TIME`.

5.6.3 Options

The simplest way to enable trace debugging for a runtime consists of setting a single option: `org.echarts.debugging`. When this option is set to `true` then the appropriate debugging and monitoring options are set so that debugger events are sent to monitors for all machine instances.

If fine-grained control over debugging is desired then there are a number of other options available similar to those for monitors (see Section 5.5.1). First, monitoring must be enabled for a runtime and for the individual machines you wish to monitor in order for debug events to be received by the monitoring subsystem. Next, it is also necessary to enable debugging for the runtime and the individual machines you wish to monitor.

A runtime's trace debugger is disabled by default. To enable it, set the `org.echarts.system.debugging` property to `true`, or invoke the `setSystemDebugging()` method. In addition to enabling debugging for the runtime, debugging must be enabled for the individual machines you want monitor. The simplest way to do this is to set the `org.echarts.machine.debugging` property to `true` or invoke the `setDefaultMachineDebugging()` method. This value is set to `false` by default. The value of this property is used as the default debugging setting for newly created machines. This default value can be overridden for a particular machine by calling the `setMachineDebugging()` method. The default behavior is that a child machine inherits the debugging setting of its parent, so if you wish to use different debugging settings in a parent and its child, you must use the `setMachineDebugging()` method.

If the system or machine debugging properties are disabled then debug events are not created and sent to monitors, regardless of a run-

time's monitor settings. Nonetheless, with monitoring enabled and debugging disabled it is possible to log non-debug events, such as `InfoEvent` instances (see Section 5.5.2).

5.6.4 Global State Output

The `local state` field described in Section 5.6.1 provides a local view of machine state. Sometimes it is also desirable to obtain a global view, that is, a snapshot of a machine's state from the root machine downwards, recursively including the state of all submachines. For even moderately complex machines, this snapshot can be hard to understand given the amount of information contained in it. Nevertheless it can sometimes be useful to periodically examine global machine state. For this reason, runtimes support printing global machine state immediately following the execution of a transition sequence (see Section 4.2). Here's an example of the output produced by the `PrettyPrintFormatter` when this option is enabled.

```
time: 2006.03.31 23:54:34:764 EST (1143867274764)
event: MachineStateEvent
root: ac2f9c:10a53cd74c2:-7fff
machine: ac2f9c:10a53cd74c2:-7fff
sequence: initializing
global state: :Example0046
    S1
        S1_1
        (S1_2)

    (S2)
```

This event corresponds to the global state of the `Example0046` machine (discussed in Section 5.6.1) immediately after the initialization sequence that causes the creation of state `S1`'s submachine. States enclosed in parentheses `()` represent non-current machine states. In this example then, the current machine state is state `S1` and its substate `S1_1`.

Root machines and their monitors are responsible for producing global state events so be sure that you minimally enable debugging for the root machine of the machine you wish to collect global state events from. Furthermore, you must explicitly inform the runtime that you wish to collect global state events since this capability is disabled by default. To enable this capability for all newly created machine instances, set the `org.echarts.machine.debugging.globalStateOutput` to `true`, or invoke the `setDefaultMachineDebuggingGlobalStateOutput()`

method. To enable or disable this capability for a particular (root) machine instance, use the `setMachineDebuggingGlobalStateOutput()` method.

5.7 Options Summary

Property	
Default Value	Section
`org.echarts.properties.dir`	
`.`	5.2
`org.echarts.system.startupMessages`	
`false`	5.3
`org.echarts.system.transitionTimerManager.class`	
`org.echarts.DefaultTransitionTimerManager`	5.4
`org.echarts.system.monitoring`	
`false`	5.5.1
`org.echarts.machine.monitoring`	
`false`	5.5.1
`org.echarts.machine.monitor.class`	
`org.echarts.monitor.PrintStreamMonitor`	5.5.1
`org.echarts.machine.monitor.filterString`	
`""`	5.5.1
`org.echarts.debugging`	
`false`	5.6.3
`org.echarts.system.debugging`	
`false`	5.6.3
`org.echarts.machine.debugging`	
`false`	5.6.3
`org.echarts.machine.debugging.globalStateOutput`	
`false`	5.6.4

6 Generating Diagrams

ECharts provides the `ech2dot` translator program for generating ECharts machine diagrams suitable for inclusion in reports, web pages, or presentations. It also provides translators for generating browsable code documentation. The documentation translators are discussed in Section 7.

All these translators share in common their dependence on the open source Graphviz `dot` program [4] for generating graphical depictions of ECharts machines. The `dot` program takes as input a text file (with a *.dot* extension) specifying the constituent nodes and arcs of a directed

graph. Its output is a graphical depiction of the graph in one of a number of possible graphical output formats e.g. PostScript, SVG, or PNG. The `.dot` files themselves are generated from ECharts machine files by the `ech2dot` translator.

Since there are so many ways to generate ECharts diagrams we've included a variety of examples here: Section 6.1 discusses generating diagrams as separate pages and Section 6.2 discusses generating diagrams suitable for embedding in other documents. In Section 6.3 we explain how one can customize a generated diagram if the default output format is unsatisfactory. Finally, in Section 6.4, we discuss how bugs in the current version of `dot` can affect a diagram's layout.

Depending on the desired diagram format, some or all of the following open source utilities (not included in the ECharts SDK) are used to assist with diagram generation: `dot`[2] for laying out and rendering the diagram in a given output format, `gs`[3] for converting PostScript to PostScript or PDF (Portable Document Format), `eps2pdf`[4] for converting encapsulated PostScript to PDF, and `inkscape`[5] for editing SVG diagrams. All examples here assume the working directory to be *runtime/java/src/examples* directory of the ECharts SDK. For information on configuring your platform to run ECharts commands like `ech2dot` see Appendix A.

6.1 Page Diagrams

Here are a number of examples of generating a diagram as a printable page. In these examples we use the following Unix conventions: the command line prompt is a percent character `%`, long commands are split over multiple lines using the backslash character `\` at the end of a line, and the command line continuation prompt for a command split over multiple lines is the 'greater than' character `>`.

Generate a 8.5 x 11 inch portrait layout diagram in PostScript format with 0.5 inch margins. We post-process the dot output with `gs` (ghostscript) in order to add explicit page size information.

```
% ech2dot --echartspath .. Example0001.ech
% dot -Tps -Gsize="7.5,10" -o Example0001.temp \
> Example0001.dot
% gs -q -dNOPAUSE -dBATCH -sPAPERSIZE=letter \
> -sDEVICE=pswrite -sOutputFile=Example0001.ps \
> Example0001.temp
```

[2] Distributed as part of the open source Graphviz package available from http://graphviz.org

[3] GNU Ghostscript is available from....

[4] Available from http://www.ctan.org/tex-archive/support/eps2pdf

[5] Available from http://inkscape.org

Generate a 8.5 x 11 inch (US letter) landscape layout diagram in PostScript format with 0.5 inch margins.

```
% ech2dot --echartspath .. Example0001.ech
% dot -Tps -Gorientation=landscape -Gsize="10,7.5" \
> -o Example0001.temp Example0001.dot
% gs -q -dNOPAUSE -dBATCH -sPAPERSIZE=letter \
> -sDEVICE=pswrite -sOutputFile=Example0001.ps \
> Example0001.temp
```

Generate a 8.5 x 11 inch landscape layout diagram in PDF (Portable Document Format) with 0.5 inch margins. We post-process the dot output with `gs` in order to translate to PDF and to add explicit page size information.

```
% ech2dot --echartspath .. Example0001.ech
% dot -Tps -Gorientation=landscape -Gsize="10,7.5" \
> -o Example0001.temp Example0001.dot
% gs -q -dNOPAUSE -dBATCH -sPAPERSIZE=letter \
> -sDEVICE=pdfwrite -sOutputFile=Example0001.pdf \
> Example0001.temp
```

Generate a 210 x 297 mm (A4) landscape layout diagram in PDF with 0.5 inch margins.

```
% ech2dot --echartspath .. Example0001.ech
% dot -Tps2 -Gsize="10.69,7.27" -Gorientation=landscape \
> -o Example0001.temp Example0001.dot
% gs -q -dNOPAUSE -dBATCH -sPAPERSIZE=a4 \
> -sDEVICE=pdfwrite -sOutputFile=Example0001.pdf \
> Example0001.temp
```

Generate a 36 x 22 inch portrait layout diagram in PostScript format with increased font size and 1 inch margins. This non-standard poster size page is suitable for printing on a 36 inch wide plotter. We post-process the dot output with `gs` in order to add explicit page size information.

```
% ech2dot --echartspath .. Example0001.ech
% dot -Tps2 -Gsize="34,20" -Efontsize=18 \
> -Elabelfontsize=12 -Gfontsize=18 \
> -o Example0001.temp Example0001.dot
% gs -q -dNOPAUSE -dBATCH -sDEVICE=pswrite \
> -dDEVICEWIDTHPOINTS=2592 -dDEVICEHEIGHTPOINTS=1584 \
> -sOutputFile=Example0001.ps Example0001.temp
```

6.2 Embedded Diagrams

Here are a few examples of how to generate a diagram for embedding in a document page, for example in a web page or a report. In these examples we use the following Unix conventions: the command line prompt is a percent character `%`, long commands are split over multiple lines using the backslash character `\` at the end of a line, and the command line continuation prompt for a command split over multiple lines is the 'greater than' character `>`.

Generate a 6 x 4 inch diagram in PNG (Portable Network Graphics) with increased font size.

```
% ech2dot --echartspath .. Example0001.ech
% dot -Tpng -Efontsize=12 -Elabelfontsize=8 \
> -Gfontsize=12 -Gsize="6,4" -o Example0001.png \
> Example0001.dot
```

Generate a 6 x 4 inch diagram in SVG (Scalable Vector Graphics) format.

```
% ech2dot --echartspath .. Example0001.ech
% dot -Tsvg -Gsize="6,4" -o Example0001.svg \
> Example0001.dot
```

Generate a 6 x 4 inch diagram in PostScript format. We post-process the `dot` output with `gs` in order to translate to EPS (encapsulated PostScript).

```
% ech2dot --echartspath .. Example0001.ech
% dot -Tps2 -Gsize="6,4" -o Example0001.temp \
> Example0001.dot
% gs -q -dNOPAUSE -dBATCH -sDEVICE=epswrite \
> -sOutputFile=Example0001.ps Example0001.temp
```

Generate a 6 x 4 inch diagram in PDF. We post-process the `gs` output with `eps2pdf` in order to translate to embeddable PDF.

```
% ech2dot --echartspath .. Example0001.ech
% dot -Tps2 -Gsize="6,4" -o Example0001.temp1 Example0001.dot
% gs -q -dNOPAUSE -dBATCH -sDEVICE=epswrite \
> -sOutputFile=Example0001.temp2 Example0001.temp1
% eps2pdf --outfile=Example0001.pdf Example0001.temp2
```

6.3 Customizing the Layout

There are four ways to customize diagrams generated with `ech2dot`:

1. editing the *.dot* file output by `ech2dot` prior to submitting it to `dot`;
2. generating an SVG format file with `dot` and then editing the diagram with `inkscape`.
3. customizing the `dot` program's output by overriding layout parameters on the `dot` command line;
4. customizing the `ech2dot` program's output by specifying custom formatters to override the default formatters;

The first option requires an understanding of the Graphviz `dot` language which we do not discuss here. Instead the interested user is referred to [5] for more information. Generating the SVG file for the second option is discussed in Section 6.1. We discuss the last two options in more detail in the following sections.

6.3.1 Overriding `dot` Layout

The `ech2dot` translator generates Graphviz *.dot* files specifying particular values for layout parameters. Of these, the most relevant for customizing the layout's appearance are its size and orientation, its label font name and sizes, its layout ratio and its printed caption.

`ech2dot` specifies the default size for a diagram to be 8 x 10.5 inches with portrait layout. The diagram size and orientation can be overridden via the `dot` command line as shown in the examples above.

The translator also specifies 8 point font size for edge (transition) and subgraph (state) labels, and the graph (machine) caption, and a 6 point font size for edge head and tail labels. The CourierNew bold font is used for all labels. As shown in the examples above, it is possible to override these font sizes on the command line. It is also possible to override the font name in a similar fashion. For example, to use Helvetica font instead of CourierNew bold:

```
% dot -Tps2 -Efontname="Helvetica" \
> -Elabelfontanme="Helvetica" -Gfontname="Helvetica" \
> -o Example0001.ps Example0001.dot
```

Note that both a machine's layout and its states' layout are dictated via `dot`'s `-G` parameter so it is not possible to specify different font names or font sizes for a machine's caption and its states.

The `ech2dot` translator also specifies the "fill" layout ratio. For this parameter value, dot attempts to fill as much of the area specified by the diagram's size as possible. It does this by expanding the size of subgraphs (states) and arcs (transitions). For simple machines with large specified size this can result in oversized nodes. To generate a diagram that satisfies the diagram's specified size without increasing

node size you can try setting the ratio to null. A diagram generated this way will have natural node sizes but is not guaranteed to fill the area specified by the diagram's size. For example:

```
% dot -Tps2 -Gratio="" -o Example0001.ps Example0001.dot
```

The default caption for an ECharts machine diagram is its (unqualified) class name. It is possible to override the caption with arbitrary text via the `dot` command line as shown in the following example:

```
% dot -Tps2 -Glabel="Arbitrary Text" -o Example0001.ps \
> Example0001.dot
```

The diagram caption can also be overridden by specifying a custom `ech2dot` formatter as explained in Section 6.3.2.

There are many other layout options available for `dot`. The interested reader is referred to [5] for more information.

6.3.2 Overriding `ech2dot` Layout

When `ech2dot` translates a *.ech* file to a *.dot* file it utilizes methods defined in modules called *formatters* to obtain formatted representations of information included in the *.dot* file. The formatted information in the *.dot* file is then included in a machine diagram output by the `dot` program. For example, a formatter can specify how a machine state or transition declaration is portrayed in a machine diagram.

The `ech2dot` program permits a user to define their own formatters if the default formatters are not suitable for their purposes. There are three formatter classes that may be overridden by the user:

label formatter This formatter defines methods for formatting state and transition labels appearing on a machine diagram.

tooltip formatter This formatter defines methods for formatting tooltips that are embedded in client-side imagemap files output by `dot`.

URL formatter This formatter defines methods for formatting URLs that are embedded in client-side imagemap, PostScript, or SVG files output by `dot`.

As described in the `ech2dot` command reference in Section 8, there are three command line options available for overriding the three formatter classes: `--label-formatter`, `--tooltip-formatter`, and `--url-formatter`, respectively. The value for each parameter must be a valid formatter (Python) class name. A formatter class is normally a subclass of the parent formatter class *translator/lib/dotmachine/DotMachineFormatter.py* included in the ECharts SDK. This class provides default definitions for methods that can be

overridden by formatter subclasses. To support the default behavior of the `ech2dot` translator, four formatter subclasses are included with the ECharts SDK in the *translator/lib/dotmachine* directory:

- *DotMachineNullFormatter.py* As its name implies, this formatter returns the empty string for all state and transition formatter method calls.
- *DotMachineCompleteFormatter.py* The methods of this formatter return the complete state or transition definition string as found in the original *.ech* file.
- *DotMachinePartialFormatter.py* The methods of this formatter return a partial state or transition definition string as found in the original *.ech* file. For example, state label strings only include the state name and the submachine class name, but any submachine parameters are excluded.
- *DotMachineCommentFormatter.py* The methods of this formatter return the user comments associated with a state or a transition.

Here are the default values of the three `ech2dot` formatters:

Formatter	Class Name
label	DotMachinePartialFormatter.DotMachinePartialFormatter
tooltip	DotMachineCommentFormatter.DotMachineCommentFormatter
URL	DotMachineFormatter.DotMachineFormatter

6.4 `dot` Layout Bugs

The open source directed graph layout program `dot` is used to produce the example machine diagrams included in this manual. The layouts produced by this program are normally excellent, however, the current version of `dot` (version 2.8) suffers from bugs that cause some transitions to be depicted incorrectly. In particular, there are two problems related to the depiction of a transition referencing the same source and target states i.e. looping transitions:

1. `dot` does not always correctly depict the machine level a transition is declared at;
2. `dot` does not always connect a transition to its referenced source or target states.

As an example of both layout problems consider the depiction of `Example0021` in Section 3.14. The graphical depiction of the transition in state `S1` shows no source state reference. A correct depiction would show the transition source connected to the border

of state `S1`. Furthermore, the transition should be depicted as extending outside `S1`'s border since the transition is declared in the `Example0018` machine, not in the state `S1` submachine. The same problems are seen in `Example0022` also in Section 3.14, `Example0025` in Section 3.15.2, `Example0027` in Section 3.16, `Example0044` in Section 3.17, `Example0029` and `Example0030` in Section 3.18, `Example0035` in Section 4.4.3, `Example0038` and `Example0039` in Section 4.6.1, `Example0040` in Section 4.6.2, and `Example0041` and `Example0042` in Section 4.7.1.

7 Generating Documentation

The `ech2doc` and `ech2javadoc` translators are used to support the generation of browsable ECharts machine code documentation. The output from both translators includes interactive SVG ECharts machine diagrams. As such, both translators rely on `ech2dot` and `dot`, the programs used to generate ECharts machine diagrams described in Section 6.

7.1 Interacting with Diagrams

There are a number of ways that a user may interact with the ECharts SVG diagrams included in the generated documentation.

1. Highlight transitions and states by rolling the cursor over them.
2. Display user comments for a machine, its transitions and its states by resting the cursor over the respective machine element.
3. Clicking on a state with a nested submachine navigates the browser to the diagram of the submachine nested in the state.
4. Pan a diagram that may be too large to fit in the viewing window. For Windows, hold the 'alt' key while dragging the diagram. For Mac OS X, hold the 'option' key while dragging.
5. Zoom in and out of a diagram. For Windows, hold the 'control' key and click to zoom in at the mouse pointer location. Hold the 'control' key and click-and-drag to select a region to zoom into. Hold the 'control' and 'shift' keys and click to zoom out. For Mac OS X, use the 'command' ('apple') key instead of the 'control' key. You can also use the zoom commands in the context menu.

Viewing and interacting with an SVG diagram requires that a web browser be configured with an SVG viewer plugin. See Section 7.4 for more information.

7.2 ech2doc

The ech2doc translator generates a web site whose pages provide a host-language independent overview of ECharts machine code. This program's output and options are similar to those provided by Java's javadoc tool. The ech2doc translator only generates documentation for ECharts .ech files and not for any supplemental host language files that may exist. Furthermore, pages generated by the ech2doc translator only include user comments for the machine, its constructors, its states and its transitions. The pages do not include user comments for host language constructs such as field or method declarations.

Here's an example of using the ech2doc translator for the examples machine package included with the ECharts SDK. In this example we assume the working directory to be *runtime/java/src/examples* of the ECharts SDK. We also adopt the following Unix conventions: the command line prompt is a percent character %, long commands are split over multiple lines using the backslash character \ at the end of a line, and the command line continuation prompt for a command split over multiple lines is the 'greater than' character >.

```
% ech2doc --echartspath .. \
> --target-directory /tmp/ech2doctest examples
```

Running this command creates a web site in the directory */tmp/ech2doctest* with the front page (*/tmp/ech2doctest/index.html*) shown in Figure 1. Clicking on the Example0001 link in the All Machines frame displays the page shown in Figure 2.

The ech2doc translator supports many of the same command line options as Java's javadoc tool. These options are discussed in the ech2doc command reference in Section 8.

7.3 ech2javadoc

Unlike the ech2doc translator discussed in Section 7.2 which is a host language independent documentation generator, the ech2javadoc generator discussed in this section is intended for ECharts machines translated to Java using ech2java. The standard program for documenting Java classes and packages is javadoc[6]. The ECharts SDK provides two utilities that enhance javadoc for documenting Java source code that includes ECharts machines: (1) the ech2javadoc translator which generates interactive SVG diagrams for ECharts machines and (2) the javadocpp program which post-processes ECharts machine class documentation produced by javadoc in order to incorporate machine diagrams into the documentation.

[6] Included with the Java SDK.

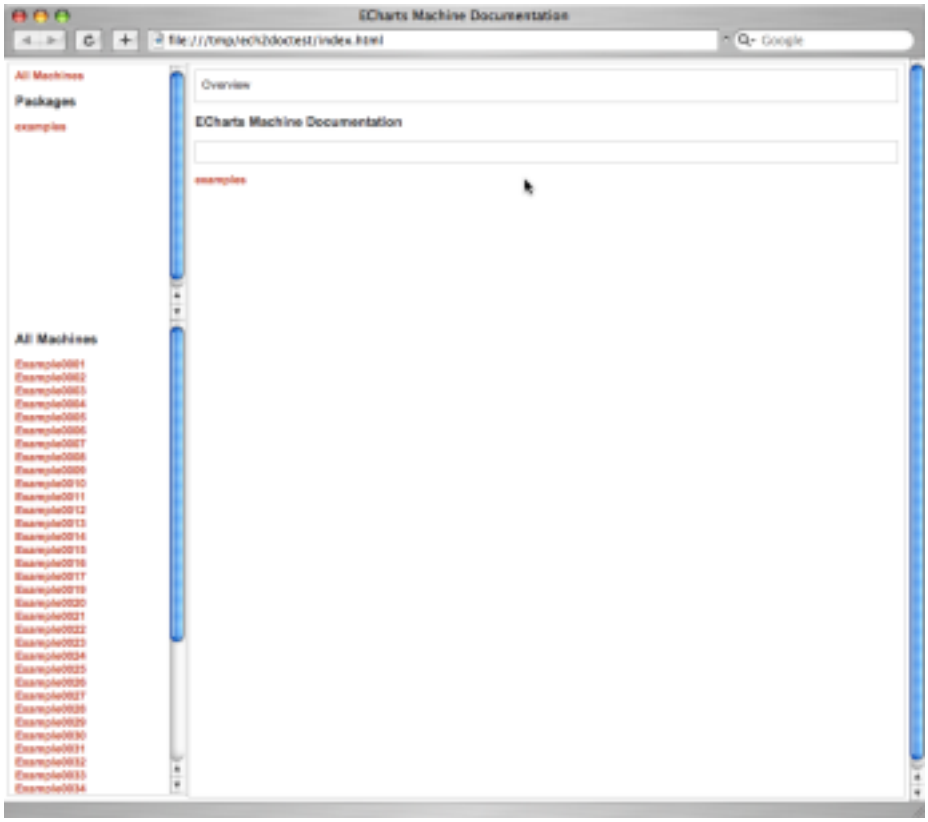

Figure 1: `ech2doc` front page

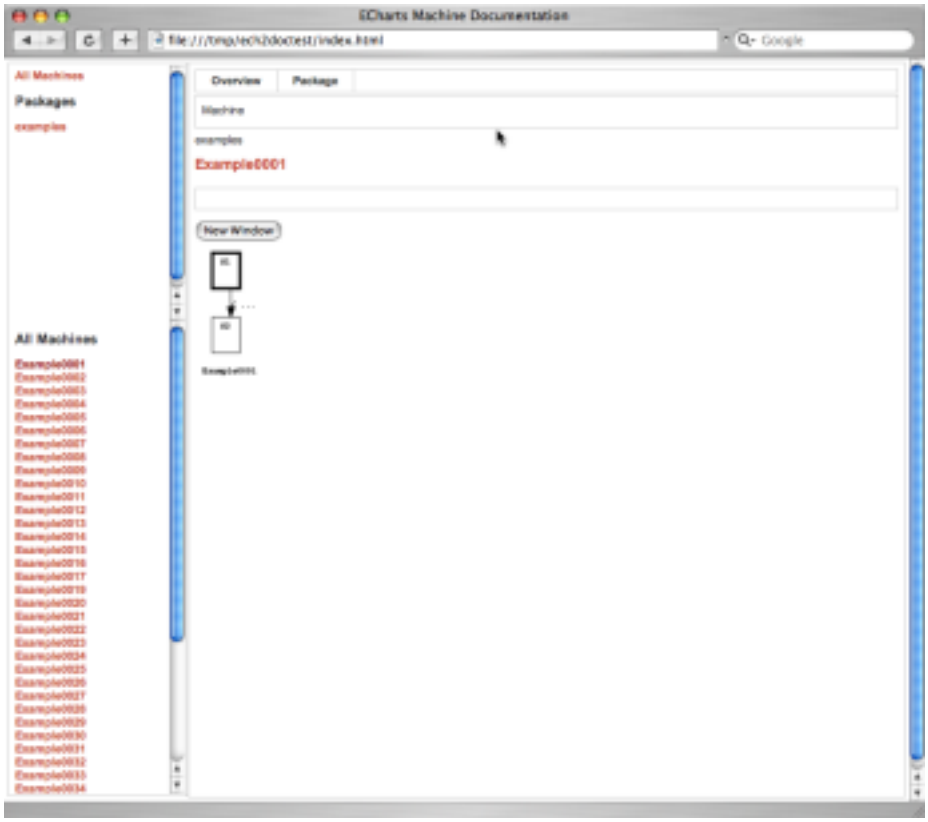

Figure 2: `ech2doc` Example0001 page

Here are the commands required to generate `javadoc` documentation for the `examples` package included with the ECharts SDK. We assume the working directory to be *runtime/java/src/examples* of the ECharts SDK.

```
% ech2javadoc --echartspath .. examples
% ech2java --echartspath .. examples
% javadoc -sourcepath .. -d /tmp/ech2javadoctest examples
% javadocpp /tmp/ech2javadoctest
```

The first command (`ech2javadoc`) translates the ECharts machines in the `examples` package into SVG diagrams and places them in the *docfiles* subdirectory of the *examples* directory in which the machines are located. The `ech2javadoc` command also creates an auxiliary `.html` file for each SVG diagram and places these files in the same *docfiles* subdirectory. Running the second command (`ech2java`) is only required if ECharts machines haven't yet been translated to Java. The ECharts machines must be translated to Java in order to be referenced by `javadoc`. The third command (`javadoc`) generates documentation for the Java files in the `examples` package and places it in the */tmp/ech2javadoctest* directory. The fourth command, `javadocpp`, post-processes the documentation generated by `javadoc` to customize the documentation for ECharts machines (see Section 8 and Appendix A for information on running this command and the ECharts translator commands). In particular, `javadocpp` removes documentation for any constructors, methods or fields that may be introduced by the `ech2java` translator but are not explicitly declared in a machine by the programmer. The program also adds a button onto documentation pages for machines. As an example, Figure 3 shows the page generated for the `Example0001` machine. Note the button labeled Machine. Clicking on this button opens a new window displaying the machine's interactive SVG diagram as shown in Figure 4.

The `ech2javadoc` translator supports a number of command line options for customizing its behavior. These are enumerated in the `ech2javadoc` command reference in Section 8.

7.4 SVG Viewers

Viewing and interacting with SVG machine diagrams embedded in documentation web pages requires that a web browser support SVG. We recommend using Adobe's SVG viewer browser plug-in[7]. Version 3.x of the plug-in is available for the Mac and Windows operating systems. We have confirmed that these plug-ins work for viewing ECharts machine diagrams with some provisos we will address momentarily. Adobe

[7] Available from `http://www.adobe.com/svg/`

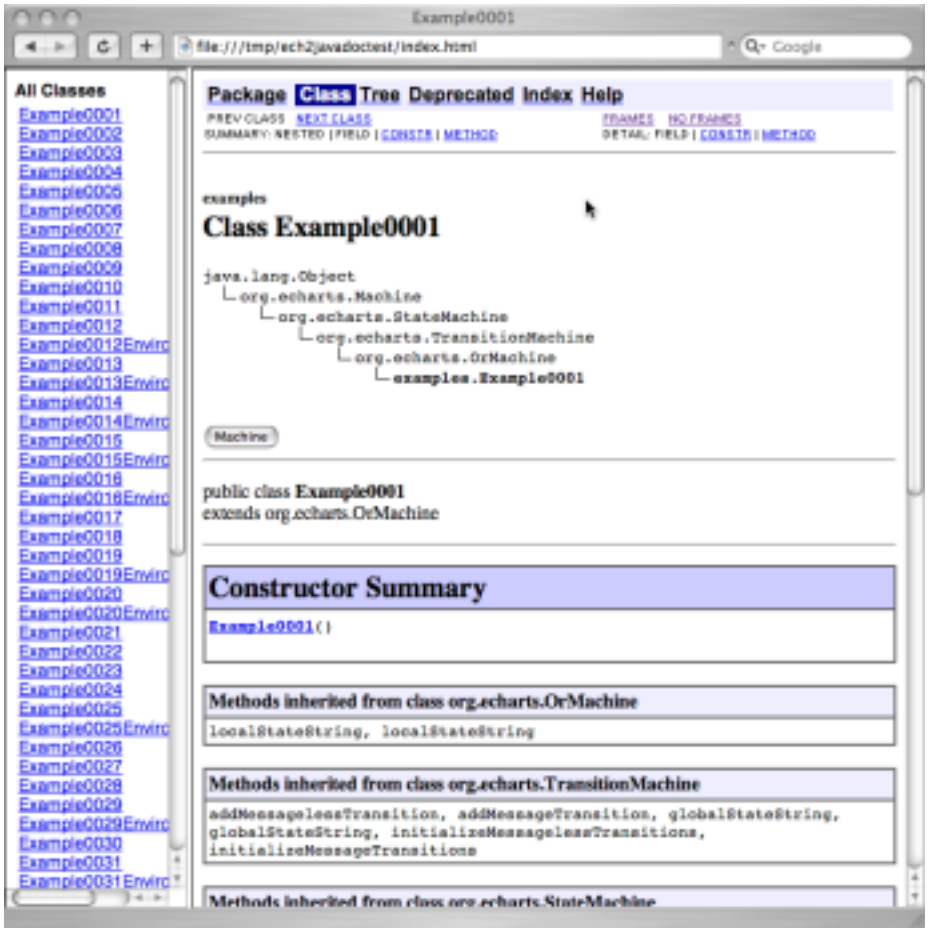

Figure 3: `ech2javadoc` Example0001 page

Figure 4: `ech2javadoc` Example0001 diagram

also provides a beta version of the plug-in for Solaris 8 and RedHat Linux 7.1-9e but we haven't tested these.

Our testing indicates that machine diagram viewing works well for the SVG plug-in/Windows/Internet Explorer combination. However, for the SVG plug-in/Mac/Safari combination we have encountered one problem. The problem relates to clicking on a state with a nested submachine in order to view the submachine's diagram. This form of user interaction does not work when the documentation web site is hosted by a web server, however it does work when the web site is accessed locally as files. This problem stems from incompatibilities between the SVG plug-in and the Safari browser and no bug fix is expected.

We should also mention that some browsers are beginning to support SVG natively. Firefox 1.5 now includes partial support for SVG. Unfortunately, attempting to view machine diagrams with Firefox reveals a number of bugs related to text rendering. Opera 9 also provides native support for SVG but it currently does not implement panning which causes large diagrams to be truncated. The nightly builds of Safari also include SVG support but we have not tested this yet.

8 Command Reference

This section of the manual is intended to serve as a quick reference to ECharts commands and their command line options. More detailed explanations of the commands can be found elsewhere in the manual.

8.1 Machine Dependencies

ECharts translators determine dependencies amongst ECharts compilation units. For example, if a machine in one *.ech* file references a machine defined in another *.ech* file, then the default behavior for most translators is to automatically translate the other file as it translates the first file. This behavior can be disabled using the `--no-dependencies` command line option. Furthermore, if the modification date of an ECharts *.ech* file is later than the modification date of its respective ECharts host language file then the *.ech* file is re-translated to ensure the translated version is up to date. These aspects of ECharts translators are similar to the support provided by Java's `javac` compiler.

8.2 ECHARTSPATH

Another similarity between `javac` and ECharts translators is the ability to define a list of directories in which to search for compilation units. While `javac` uses the `CLASSPATH` environment variable, ECharts uses the `ECHARTSPATH` environment variable. When the ECharts translator needs to locate a compilation unit upon which another compilation unit depends, it searches each directory listed in the `ECHARTSPATH` variable in the order they are listed until it locates the compilation unit. Similar to `javac`, ECharts translators reference the `ECHARTSPATH` environment variable or the `--echartspath` command-line argument. For example, in `Example0001` (see Section 2.1), the `ech2java` translator is invoked as follows:

```
% ech2java --echartspath .. Example0001.ech
```

However, an equivalent way to invoke the translator would be to initialize the `ECHARTSPATH` environment variable. Here we assume a Unix environment using the bash shell.

```
% export ECHARTSPATH=..; ech2java Example0001.ech
```

If the `ECHARTSPATH` environment variable is defined and an ECharts translator is invoked with the `--echartspath` command-line option, then the value defined by the command-line option will override the value defined by the environment variable.

8.3 ech2java

The `ech2java` translator translates ECharts *.ech* machine definition files to Java *.java* files. The `ech2java` translator is introduced in Section 2.1.

Usage:

```
ech2java [options] [echartsfilenames] [echartspackagenames]
```

Summary:

Translate specified `echartsfilenames` and `echartspackagenames` to Java *.java* files. For each specified package, translate all *.ech* files in that package.

Options:

--echartspath ':' separated list (';' on Windows) specifying directories to search for the specified `echartsfilenames` and `echartspackagenames`. The default value is the current working directory.

--no-dependencies Forces translator to not translate *.ech* files upon which the specified `echartsfilenames` and `echartspackagenames` depend. The default behavior is to translate dependencies.

--target-directory Directory path specifying where to write translated files to. If subdirectories on path do not exist then they are created. The default value is the source directory of the specified `echartsfilenames` and `echartspackagenames`.

--package-subdirectory Package subdirectory name relative to the target directory specifying where to write translated files to. The default value is the empty string.

--subpackages ':' separated list (';' on Windows) list specifying package names recursively searched for *.ech* files. This is similar to the `echartspackagenames` command line argument except that package subdirectories are searched recursively for this option. The default value is the empty string.

--version Print translator version and exit.

--help Print help message and exit.

8.4 ech2dot

The `ech2dot` translator translates ECharts *.ech* machine definition files to Graphviz *.dot* files. The `ech2dot` translator is discussed in detail in Section 6.

Usage:

```
ech2dot [options] [echartsfilenames] [echartspackagenames]
```

Summary:

Translate specified `echartsfilenames` and `echartspackagenames` to Graphviz *.dot* files. For each specified package, translate all *.ech* files in that package.

Options:

--echartspath ':' separated list (';' on Windows) specifying directories to search for the specified `echartsfilenames` and `echartspackagenames`. The default value is the current working directory.

--no-dependencies Forces translator to not translate *.ech* files upon which the specified `echartsfilenames` and `echartspackagenames` depend. The default behavior is to translate dependencies.

--target-directory Directory path specifying where to write translated files to. If subdirectories on path do not exist then they are created. The default value is the source directory of the specified `echartsfilenames` and `echartspackagenames`.

--package-subdirectory Package subdirectory name relative to the target directory specifying where to write translated files to. The default value is the empty string.

--subpackages ':' separated list (';' on Windows) list specifying package names recursively searched for *.ech* files. This is similar to the `echartspackagenames` command line argument except that package subdirectories are searched recursively for this option. The default value is the empty string.

--label-formatter Fully qualified Python class name of label formatter. Default value is DotMachinePartialFormatter.DotMachinePartialFormatter.

--tooltip-formatter Fully qualified Python class name of tooltip formatter. Default value is DotMachineCommentFormatter.DotMachineCommentFormatter.

--url-formatter Fully qualified Python class name of URL formatter. Default value is DotMachineFormatter.DotMachineFormatter.

--version Print translator version and exit.

--help Print help message and exit.

8.5 `ech2doc`

The `ech2doc` translator translates ECharts *.ech* machine definition files to *.html* and *.svg* files. The `ech2doc` translator is discussed in detail in Section 7.2.

Usage:

`ech2doc [options] [echartsfilenames] [echartspackagenames]`

Summary:

Translate specified `echartsfilenames` and `echartspackagenames` to *.html* and *.svg* files. For each specified package, translate all *.ech* files in that package. Note that the Graphviz `dot` program must be installed to run this translator.

Options:

--echartspath ':' separated list (';' on Windows) specifying directories to search for the specified `echartsfilenames` and `echartspackagenames`. The default value is the current working directory.

--no-dependencies Forces translator to not translate *.ech* files upon which the specified `echartsfilenames` and `echartspackagenames` depend. The default behavior is to translate dependencies.

--target-directory Directory path specifying where to write translated files to. If subdirectories on path do not exist then they are created. The default value is the source directory of the specified `echartsfilenames` and `echartspackagenames`.

--subpackages ':' separated list (';' on Windows) list specifying package names recursively searched for *.ech* files. This is similar to the `echartspackagenames` command line argument except that package subdirectories are searched recursively for this option. The default value is the empty string.

--dot-path 'dot' program file path. Default value is 'dot'.

--document-title Documentation title HTML. Default value is 'ECharts Machine Documentation'.

--window-title Browser window title. Default value is document title.

--overview Path to documentation overview file shown on first page of documentation. Overview may utilize HTML tags. Default is no overview file.

--header Documentation header HTML. Default value is empty string.

--header Documentation footer HTML. Default value is empty string.

--svg-only Only generate *.html* and *.svg* files to support browsable SVG machine diagrams. Does not generate *.html* files for navigating machine packages or classes. Default behavior is to generate all *.html* files.

--svg-fontsize Font size in points for text associated with SVG machine diagrams. Default value is '8'.

--version Print translator version and exit.

--help Print help message and exit.

8.6 ech2javadoc

The `ech2javadoc` translator translates ECharts *.ech* machine definition files to *.html* and *.svg* files. The `ech2javadoc` translator is discussed in detail in Section 7.3.

Usage:

```
ech2javadoc [options] [echartsfilenames] [echartspackagenames]
```

Summary:

Translate specified `echartsfilenames` and `echartspackagenames` to *.html* and *.svg* files. For each specified package, translate all *.ech* files in that package. The target directory for translated files is a directory named *doc-files* that is a subdirectory of the *.ech* file source directory. Note that the Graphviz `dot` program must be installed to run this translator.

Options:

--echartspath ':' separated list (';' on Windows) specifying directories to search for the specified `echartsfilenames` and `echartspackagenames`. The default value is the current working directory.

--target-directory Directory path specifying where to write translated files to. If subdirectories on path do not exist then they are created. The default value is the source directory of the specified `echartsfilenames` and `echartspackagenames`.

--subpackages ':' separated list (';' on Windows) list specifying package names recursively searched for *.ech* files. This is similar to the `echartspackagenames` command line argument except that package subdirectories are searched recursively for this option. The default value is the empty string.

--dot-path 'dot' program file path. Default value is 'dot'.

--version Print translator version and exit.

--help Print help message and exit.

8.7 javadocpp

The `javadocpp` program post-processes HTML documentation generated by `javadoc` to provide customized views of ECharts machine class documentation. The `javadocpp` program is introduced in Section 7.3.

Usage:

```
javadocpp [options] [javadocdirs]
```

Summary:

Post-process specified `javadocdirs`. Each specified directory should contain `javadoc` HTML documentation. For each specified directory, post-process all HTML files corresponding to ECharts machine classes and recursively invoke the program for any subdirectories.

Options:

--version Print program version and exit.

--help Print help message and exit.

9 Coming Soon!

A number of features are currently under development for release with the next ECharts version. These features are required for a number of research projects involving ECharts at AT&T, however, we expect they will be of general use too.

9.1 Unit and Regression Testing

The *javamachine* runtime currently includes a rudimentary unit test and regression test framework. These are used to support the *javamachine* runtime test suite found in the *runtime/java/src/test* directory of the ECharts SDK. The next version of ECharts will provide enhanced support for this framework.

9.2 Indexed Machine Array Access

Support for referencing a machine array element (see Section 3.14) at a specific integer index will be supported. This includes specifying the index for a newly created machine array element.

9.3 Internal Ports

Internal ports are intended to make communications between a machine, its ancestors and its descendants explicit. Messages are sent and received on internal ports using the same syntax currently used to send and receive messages on external ports (see Section 3.13). The priority of an internal port will be between that of a timed transition port and an external port (see Section 4.3).

9.4 Past and Future State References

Past and future state references permit the state of a submachine to be referenced by a transition when the submachine is not part of the overall machine's current state. This facility generalizes the UML Statecharts notion of a SynchState [6] in a semantically appealing way.

9.5 Exception Handling

The current version of ECharts requires programmers to handle run-time exceptions using the exception handling facility provided by the underlying host language. The next version will support exception handling in the ECharts language itself.

10 Coming Later!

We've already devoted a lot of time to the features described in this section, however, their relative priority for our own work has meant that their implementation must wait until the features described in the previous section are implemented.

10.1 Machine Inheritance

Conspicuously absent from the current version of ECharts is support for machine inheritance. Our approach to inheritance will permit sub-machines to add or override states and transitions in machine super-classes.

10.2 Machine Variables

The adventurous ECharts programmer may have noticed that the current version of ECharts supports recursive machine invocation, that is to say, the ability for a machine to directly or indirectly invoke itself. With appropriately specified constructors and transition guards, one can even avoid infinite recursion. However this is more of a party trick than a fundamentally useful feature. What ECharts requires to make it a full-fledged (Turing complete) language is the ability to perform operations directly on a machine at runtime. Our approach will permit machines to be assigned to machine variables and machine variables to be assigned to machine states. Our goal in providing this support is not to claim Turing completeness; rather, it is to simplify programming tasks we have encountered.

References

[1] Gregory W. Bond. Timed transition activation semantics in Statecharts. Technical Report TD-6H4L5S, AT&T, 2005. Available from: http://echarts.org. 47

[2] Gregory W. Bond and Healfdene H. Goguen. ECharts: balancing design and implementation. In M.H. Hamza, editor, *Proceedings of the 6th IASTED International Conference on Software Engineering*

and Applications (SEA 2002), pages 149–155. ACTA Press, 2002. Available from: http://echarts.org. 4

[3] Gregory W. Bond and Healfdene H. Goguen. ECharts: From lab to production. Technical Report TD-6FSMWT, AT&T, 2005. Available from: http://echarts.org. 5, 63

[4] Emden R. Gansner and Stephen C. North. An open graph visualization system and its applications to software engineering. *Software — Practice and Experience*, 30(11):1203–1233, 2000. Available from: http://graphviz.org. 88

[5] Eleftherios Koutsofios and Stephen C. North. *Drawing graphs with dot*. Available from http://graphviz.org. 92, 93

[6] Object Management Group. *OMG Unified Modeling Language Superstructure Specification, version 2.0*. Object Management Group, August 2005. Available from: http://www.omg.org. 4, 107

A Building and Using ECharts

The ECharts SDK is supported on Solaris, Mac OS X, Linux, Windows and Eclipse. A number of additional open source software packages are required to build and use ECharts. We enumerate these packages here and provide instructions on how to configure your platform to use ECharts with these packages.

A.1 Software Requirements

The required software for building and using the ECharts SDK are Python 2.2 or greater and the Java 2 Platform, Standard Edition (JDK 1.4.2 or greater) and a suitable build tool. Python is bundled with the Linux and Mac OS X platforms. On Solaris and Windows a separate download is required from http://python.org. Java is bundled only with Mac OS X. Other platforms require a separate download from http://java.sun.com.

There are three possible build tools for building the ECharts SDK: the Eclipse SDK 3.1.2 or greater, Apache Ant 1.6.3 or greater, or GNU Make and the GNU BASH shell. Eclipse is available from http://www.eclipse.org for all platforms. Apache Ant is available from http://ant.apache.org for all platforms. GNU Make and BASH are bundled with Linux and Mac OS X (GNU Make is included with the Mac OS X Xcode Developer Tools). A separate download is required for Solaris from http://gnu.org. We do not recommend using Make or BASH on Windows since we have not tested them with the ECharts build. Instead, we recommend the use of Eclipse or Ant.

For generating machine diagrams or machine documentation (see Sections 6 and 7), Graphviz version 1.13 or greater is required. This

is available as a separate download from http://graphviz.org for all platforms.

A.2 Building ECharts

The top-level build builds the ECharts language parser and the ECharts Java runtime library `echarts.jar`. This is a pre-requisite to translating ECharts machines, as described in Appendix A.3.

Before kicking off the build, ensure your platform's `PATH` environment variable references the Python executable directory. When building with Apache Ant, simply execute `ant all` in the ECharts top-level directory. When using Eclipse, select the 'all' target of the top-level *build.xml* file by invoking the file's 'Run As -> Ant Build...' contextual menu option. Note that Eclipse should be configured to use the external Java 2 JDK, not Ant's internal Java compiler. This is accomplished via the 'Java -> Installed JREs' node of the Eclipse 'Preferences' panel. When building with GNU Make, the value of the `JDK_HOME` environment variable must be set to the location of your JDK installation prior to invoking Make. One way to do this is from the BASH shell is as follows. Here we assume the JDK is installed in */usr/java/j2sdk1.4.2_10*.

```
% JDK_HOME=/usr/java/j2sdk1.4.2_10 make all
```

A.3 Using ECharts

Once the top-level build has been performed (see Appendix A.2) you are ready to start using the ECharts SDK to translate ECharts machines. The ECharts translators (e.g. `ech2java`) and associated programs (e.g. `javadocpp`) are Python scripts. They can be invoked either directly or indirectly.

A.3.1 Direct Invocation

On Unix platforms (Mac OS X, Solaris and Linux), the ECharts translators and programs can be invoked from the command line. To invoke a translator or program without having to prefix it with its path, the path to the translator executables can be added to the system path as follows. Here we assume the path to the ECharts SDK is *ECHARTS_HOME*, and that we are using the BASH shell.

```
export PATH=$PATH:ECHARTS_HOME/translator/bin
```

On Windows, you can invoke the ECharts translators and programs from the Windows `Cmd.exe` command line shell. However, invocations must be prefixed with an explicit invocation of Python. Here's an example of invoking `ech2java` where we assume the path to the ECharts SDK is *ECHARTS_HOME*.

```
C:\> python ECHARTS_HOME\translator\bin\ech2java --version
```

A.3.2 Indirect Invocation

Translators and programs may also be invoked indirectly via Apache Ant on all platforms, or via GNU Make on Unix platforms (Mac OS X, Solaris and Linux). See the Ant and Make build files located in *runtime/java/src/examples* of the ECharts SDK for an example of invoking `ech2java`. Also see the Make and Ant build files located in *runtime/java/doc* in the ECharts SDK for an example of invoking `ech2javadoc` and `javadocpp`.

B Licenses

All software in the ECharts SDK is open source software.

Unless otherwise indicated the software in the ECharts SDK is part of the ECharts package copyright (c) 2006 AT&T Corp. and is licensed under the Common Public License, Version 1.0 by AT&T Corp. (see http://www.opensource.org/licenses/cpl1.0.php).

The other license applicable to this SDK is the ANTLR 2 License (see http://www.antlr.org/license.html). This license covers the parser generator underlying ECharts machine translators.

B.1 Common Public License v1.0

THE ACCOMPANYING PROGRAM IS PROVIDED UNDER THE TERMS OF THIS COMMON PUBLIC LICENSE ("AGREEMENT"). ANY USE, REPRODUCTION OR DISTRIBUTION OF THE PROGRAM CONSTITUTES RECIPIENT'S ACCEPTANCE OF THIS AGREEMENT.

1. DEFINITIONS

   "Contribution" means:

   (a) in the case of the initial Contributor, the initial code and documentation distributed under this Agreement, and

   (b) in the case of each subsequent Contributor:

      i. changes to the Program, and

      ii. additions to the Program;

   where such changes and/or additions to the Program originate from and are distributed by that particular Contributor. A Contribution 'originates' from a Contributor if it was added to the Program by such Contributor itself or anyone acting on such Contributor's behalf. Contributions do not include additions to the Program which: (i) are separate modules of software distributed

in conjunction with the Program under their own license agreement, and (ii) are not derivative works of the Program.

"Contributor" means any person or entity that distributes the Program.

"Licensed Patents " mean patent claims licensable by a Contributor which are necessarily infringed by the use or sale of its Contribution alone or when combined with the Program.

"Program" means the Contributions distributed in accordance with this Agreement.

"Recipient" means anyone who receives the Program under this Agreement, including all Contributors.

2. GRANT OF RIGHTS

   (a) Subject to the terms of this Agreement, each Contributor hereby grants Recipient a non-exclusive, worldwide, royalty-free copyright license to reproduce, prepare derivative works of, publicly display, publicly perform, distribute and sublicense the Contribution of such Contributor, if any, and such derivative works, in source code and object code form.

   (b) Subject to the terms of this Agreement, each Contributor hereby grants Recipient a non-exclusive, worldwide, royalty-free patent license under Licensed Patents to make, use, sell, offer to sell, import and otherwise transfer the Contribution of such Contributor, if any, in source code and object code form. This patent license shall apply to the combination of the Contribution and the Program if, at the time the Contribution is added by the Contributor, such addition of the Contribution causes such combination to be covered by the Licensed Patents. The patent license shall not apply to any other combinations which include the Contribution. No hardware per se is licensed hereunder.

   (c) Recipient understands that although each Contributor grants the licenses to its Contributions set forth herein, no assurances are provided by any Contributor that the Program does not infringe the patent or other intellectual property rights of any other entity. Each Contributor disclaims any liability to Recipient for claims brought by any other entity based on infringement of intellectual property rights or otherwise. As a condition to exercising the rights and licenses granted hereunder, each Recipient hereby assumes sole responsibility to secure any other intellectual property rights needed, if any. For example, if a third party patent license is required to allow Recipient to distribute the Program, it is Recipient's responsibility to acquire that license before distributing the Program.

(d) Each Contributor represents that to its knowledge it has sufficient copyright rights in its Contribution, if any, to grant the copyright license set forth in this Agreement.

3. REQUIREMENTS

A Contributor may choose to distribute the Program in object code form under its own license agreement, provided that:

(a) it complies with the terms and conditions of this Agreement; and

(b) its license agreement:

i. effectively disclaims on behalf of all Contributors all warranties and conditions, express and implied, including warranties or conditions of title and non-infringement, and implied warranties or conditions of merchantability and fitness for a particular purpose;

ii. effectively excludes on behalf of all Contributors all liability for damages, including direct, indirect, special, incidental and consequential damages, such as lost profits;

iii. states that any provisions which differ from this Agreement are offered by that Contributor alone and not by any other party; and

iv. states that source code for the Program is available from such Contributor, and informs licensees how to obtain it in a reasonable manner on or through a medium customarily used for software exchange.

When the Program is made available in source code form:

(a) it must be made available under this Agreement; and

(b) a copy of this Agreement must be included with each copy of the Program.

Contributors may not remove or alter any copyright notices contained within the Program.

Each Contributor must identify itself as the originator of its Contribution, if any, in a manner that reasonably allows subsequent Recipients to identify the originator of the Contribution.

4. COMMERCIAL DISTRIBUTION

Commercial distributors of software may accept certain responsibilities with respect to end users, business partners and the like. While this license is intended to facilitate the commercial use of the Program, the Contributor who includes the Program in a commercial product offering should do so in a manner which does not create potential liability for other Contributors. Therefore,

if a Contributor includes the Program in a commercial product offering, such Contributor ("Commercial Contributor") hereby agrees to defend and indemnify every other Contributor ("Indemnified Contributor") against any losses, damages and costs (collectively "Losses") arising from claims, lawsuits and other legal actions brought by a third party against the Indemnified Contributor to the extent caused by the acts or omissions of such Commercial Contributor in connection with its distribution of the Program in a commercial product offering. The obligations in this section do not apply to any claims or Losses relating to any actual or alleged intellectual property infringement. In order to qualify, an Indemnified Contributor must: a) promptly notify the Commercial Contributor in writing of such claim, and b) allow the Commercial Contributor to control, and cooperate with the Commercial Contributor in, the defense and any related settlement negotiations. The Indemnified Contributor may participate in any such claim at its own expense.

For example, a Contributor might include the Program in a commercial product offering, Product X. That Contributor is then a Commercial Contributor. If that Commercial Contributor then makes performance claims, or offers warranties related to Product X, those performance claims and warranties are such Commercial Contributor's responsibility alone. Under this section, the Commercial Contributor would have to defend claims against the other Contributors related to those performance claims and warranties, and if a court requires any other Contributor to pay any damages as a result, the Commercial Contributor must pay those damages.

5. NO WARRANTY

   EXCEPT AS EXPRESSLY SET FORTH IN THIS AGREEMENT, THE PROGRAM IS PROVIDED ON AN "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, EITHER EXPRESS OR IMPLIED INCLUDING, WITHOUT LIMITATION, ANY WARRANTIES OR CONDITIONS OF TITLE, NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Each Recipient is solely responsible for determining the appropriateness of using and distributing the Program and assumes all risks associated with its exercise of rights under this Agreement, including but not limited to the risks and costs of program errors, compliance with applicable laws, damage to or loss of data, programs or equipment, and unavailability or interruption of operations.

6. DISCLAIMER OF LIABILITY

   EXCEPT AS EXPRESSLY SET FORTH IN THIS AGREE-

MENT, NEITHER RECIPIENT NOR ANY CONTRIBUTORS SHALL HAVE ANY LIABILITY FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING WITHOUT LIMITATION LOST PROFITS), HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OR DISTRIBUTION OF THE PROGRAM OR THE EXERCISE OF ANY RIGHTS GRANTED HEREUNDER, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

7. GENERAL

   If any provision of this Agreement is invalid or unenforceable under applicable law, it shall not affect the validity or enforceability of the remainder of the terms of this Agreement, and without further action by the parties hereto, such provision shall be reformed to the minimum extent necessary to make such provision valid and enforceable.

   If Recipient institutes patent litigation against a Contributor with respect to a patent applicable to software (including a cross-claim or counterclaim in a lawsuit), then any patent licenses granted by that Contributor to such Recipient under this Agreement shall terminate as of the date such litigation is filed. In addition, if Recipient institutes patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Program itself (excluding combinations of the Program with other software or hardware) infringes such Recipient's patent(s), then such Recipient's rights granted under Section 2(b) shall terminate as of the date such litigation is filed.

   All Recipient's rights under this Agreement shall terminate if it fails to comply with any of the material terms or conditions of this Agreement and does not cure such failure in a reasonable period of time after becoming aware of such noncompliance. If all Recipient's rights under this Agreement terminate, Recipient agrees to cease use and distribution of the Program as soon as reasonably practicable. However, Recipient's obligations under this Agreement and any licenses granted by Recipient relating to the Program shall continue and survive.

   Everyone is permitted to copy and distribute copies of this Agreement, but in order to avoid inconsistency the Agreement is copyrighted and may only be modified in the following manner. The Agreement Steward reserves the right to publish new versions (including revisions) of this Agreement from time to time. No one other than the Agreement Steward has the right to modify

this Agreement. IBM is the initial Agreement Steward. IBM may assign the responsibility to serve as the Agreement Steward to a suitable separate entity. Each new version of the Agreement will be given a distinguishing version number. The Program (including Contributions) may always be distributed subject to the version of the Agreement under which it was received. In addition, after a new version of the Agreement is published, Contributor may elect to distribute the Program (including its Contributions) under the new version. Except as expressly stated in Sections 2(a) and 2(b) above, Recipient receives no rights or licenses to the intellectual property of any Contributor under this Agreement, whether expressly, by implication, estoppel or otherwise. All rights in the Program not expressly granted under this Agreement are reserved.

This Agreement is governed by the laws of the State of New York and the intellectual property laws of the United States of America. No party to this Agreement will bring a legal action under this Agreement more than one year after the cause of action arose. Each party waives its rights to a jury trial in any resulting litigation.

B.2 ANTLR 2 License

We reserve no legal rights to the ANTLR–it is fully in the public domain. An individual or company may do whatever they wish with source code distributed with ANTLR or the code generated by ANTLR, including the incorporation of ANTLR, or its output, into commerical software.

We encourage users to develop software with ANTLR. However, we do ask that credit is given to us for developing ANTLR. By "credit", we mean that if you use ANTLR or incorporate any source code into one of your programs (commercial product, research project, or otherwise) that you acknowledge this fact somewhere in the documentation, research report, etc... If you like ANTLR and have developed a nice tool with the output, please mention that you developed it using ANTLR. In addition, we ask that the headers remain intact in our source code. As long as these guidelines are kept, we expect to continue enhancing this system and expect to make other tools available as they are completed.